UGC-NET

CHAPTERWISE QUESTION BANK

COMPUTER SCIENCE & APPLICATIONS

- Full Explanation for Every Question
- Strictly Based on Latest Syllabus
- Matches the Real Exam Pattern Exactly
- Predicted Questions for Upcoming Exam

NODIA

Scan QR for Demo PDF

Nodia

GET ALL STUDY MATERIAL IN ONE APP!

Textbooks

Previous Year
Chapterwise Question Banks

30 Sets
Sample Papers

Worksheets

FREE!

DOWNLOAD NOW!

SCAN & DOWNLOAD

GET IT ON Google Play

UGC NET 2026

COMPUTER SCIENCE & APPLICATIONS

CHAPTER-WISE MCQ QUESTION BANK

INCLUDING SOURCE-BASED QUESTIONS

NODIA

Downloaded from www.pdf.tube

UGC NET Computer Science and Applications
Edition February 2026
Published by :NODIA
An Imprint of NODIA EDITIONS LLP

NODIA

Information contained in this book has been obtained by author, from sources believes to be reliable. However, neither NODIA nor its author guarantee the accuracy or completeness of any information herein and NODIA nor its author shall be responsible for any error, omissions, or damages arising out of use of this information. This book is published with the understanding that NODIA and its author are supplying information but are not attempting to render engineering or other professional services.

© NODIA EDITIONS LLP

MRP Rs 710.00

available at
amazon

This book is available on amazon and flipkart only and not available in market.

NODIA EDITIONS LLP
125, Sector 6, Vidyadhar Nagar, Jaipur 302039
Phone :+91 9024037387
email : mail@nodia.in

CONTENTS

Chap 01.	Discrete Structures and Optimization	13-76
Chap 02.	Computer System Architecture	77-148
Chap 03.	Programming Languages and Computer Graphics	149-220
Chap 04.	Database Management Systems	221-291
Chap 05.	System Software and Operating System	292-360
Chap 06.	Software Engineering	361-430
Chap 07.	Data Structures and Algorithms	431-499
Chap 08.	Theory of Computation and Compilers	500-565
Chap 09.	Data Communication and Computer Networks	566-633
Chap 10.	Artificial Intelligence (AI)	634-703

****

Click Here to Buy book on Amazon

UNIVERSITY GRANTS COMMISSION
NET BUREAU

NET SYLLABUS

Subject : COMPUTER SCIENCE AND APPLICATIONS

Code No.:(87)

Unit - 1 : Discrete Structures and Optimization

Mathematical Logic: Propositional and Predicate Logic, Propositional Equivalences, Normal Forms, Predicates and Quantifiers, Nested Quantifiers, Rules of Inference.

Sets and Relations: Set Operations, Representation and Properties of Relations, Equivalence Relations, Partially Ordering.

Counting, Mathematical Induction and Discrete Probability: Basics of Counting, Pigeonhole Principle, Permutations and Combinations, Inclusion- Exclusion Principle, Mathematical Induction, Probability, Bayes' Theorem.

Group Theory: Groups, Subgroups, Semi Groups, Product and Quotients of Algebraic Structures, Isomorphism, Homomorphism, Automorphism, Rings, Integral Domains, Fields, Applications of Group Theory.

Graph Theory: Simple Graph, Multigraph, Weighted Graph, Paths and Circuits, Shortest Paths in Weighted Graphs, Eulerian Paths and Circuits, Hamiltonian Paths and Circuits, Planner graph, Graph Coloring, Bipartite Graphs, Trees and Rooted Trees, Prefix Codes, Tree Traversals, Spanning Trees and Cut-Sets.

Boolean Algebra: Boolean Functions and its Representation, Simplifications of Boolean Functions.

Optimization: Linear Programming - Mathematical Model, Graphical Solution, Simplex and Dual Simplex Method, Sensitive Analysis; Integer Programming, Transportation and Assignment Models, PERT-CPM: Diagram Representation, Critical Path Calculations, Resource Levelling, Cost Consideration in Project Scheduling.

Unit - 2 : Computer System Architecture

NODA

Digital Logic Circuits and Components: Digital Computers, Logic Gates, Boolean Algebra, Map Simplifications, Combinational Circuits, Flip-Flops, Sequential Circuits, Integrated Circuits, Decoders, Multiplexers, Registers and Counters, Memory Unit.

Data Representation: Data Types, Number Systems and Conversion, Complements, Fixed Point Representation, Floating Point Representation, Error Detection Codes, Computer Arithmetic - Addition, Subtraction, Multiplication and Division Algorithms.

Register Transfer and Microoperations: Register Transfer Language, Bus and Memory Transfers, Arithmetic, Logic and Shift Microoperations.

Basic Computer Organization and Design: Stored Program Organization and Instruction Codes, Computer Registers, Computer Instructions, Timing and Control, Instruction Cycle, Memory-Reference Instructions, Input-Output, Interrupt.

Programming the Basic Computer: Machine Language, Assembly Language, Assembler, Program Loops, Subroutines, Input-Output Programming.

Microprogrammed Control: Control Memory, Address Sequencing, Design of Control Unit.

Central Processing Unit: General Register Organization, Stack Organization, Instruction Formats, Addressing Modes, RISC Computer, CISC Computer.

Pipeline and Vector Processing: Parallel Processing, Pipelining, Arithmetic Pipeline, Instruction Pipeline, Vector Processing Array Processors.

Input-Output Organization: Peripheral Devices, Input-Output Interface, Asynchronous Data Transfer, Modes of Transfer, Priority Interrupt, DMA, Serial Communication.

Memory Hierarchy: Main Memory, Auxillary Memory, Associative Memory, Cache Memory, Virtual Memory, Memory Management Hardware.

Multiprocessors: Characteristics of Multiprocessors, Interconnection Structures, Interprocessor Arbitration, Interprocessor Communication and Synchronization, Cache Coherence, Multicore Processors.

Unit - 3 : Programming Languages and Computer Graphics

Language Design and Translation Issues: Programming Language Concepts, Paradigms and Models, Programming Environments, Virtual Computers and Binding Times, Programming Language Syntax, Stages in Translation, Formal Transition Models.

Elementary Data Types: Properties of Types and Objects; Scalar and Composite Data Types.

Programming in C: Tokens, Identifiers, Data Types, Sequence Control, Subprogram Control, Arrays, Structures, Union, String, Pointers, Functions, File Handling, Command Line Argumaents, Preprocessors.

Downloaded from www.pdf.tube

Object Oriented Programming: Class, Object, Instantiation, Inheritance, Encapsulation, Abstract Class, Polymorphism.

Programming in C++: Tokens, Identifiers, Variables and Constants; Data types, Operators, Control statements, Functions Parameter Passing, Virtual Functions, Class and Objects; Constructors and Destructors; Overloading, Inheritance, Templates, Exception and Event Handling; Streams and Files; Multifile Programs.

Web Programming: HTML, DHTML, XML, Scripting, Java, Servlets, Applets.

Computer Graphics: Video-Display Devices, Raster-Scan and Random-Scan Systems; Graphics Monitors, Input Devices, Points and Lines; Line Drawing Algorithms, Mid-Point Circle and Ellipse Algorithms; Scan Line Polygon Fill Algorithm, Boundary-Fill and Flood-Fill.

2-D Geometrical Transforms and Viewing: Translation, Scaling, Rotation, Reflection and Shear Transformations; Matrix Representations and Homogeneous Coordinates; Composite Transforms, Transformations Between Coordinate Systems, Viewing Pipeline, Viewing Coordinate Reference Frame, Window to View-Port Coordinate Transformation, Viewing Functions, Line and Polygon Clipping Algorithms.

3-D Object Representation, Geometric Transformations and Viewing: Polygon Surfaces, Quadric Surfaces, Spline Representation, Bezier and B-Spline Curves; Bezier and B-Spline Surfaces; Illumination Models, Polygon Rendering Methods, Viewing Pipeline and Coordinates; General Projection Transforms and Cipping.

Unit – 4 : Database Management Systems

NODA

Database System Concepts and Architecture: Data Models, Schemas, and Instances; Three-Schema Architecture and Data Independence; Database Languages and Interfaces; Centralized and Client/Server Architectures for DBMS.

Data Modeling: Entity-Relationship Diagram, Relational Model - Constraints, Languages, Design, and Programming, Relational Database Schemas, Update Operations and Dealing with Constraint Violations; Relational Algebra and Relational Calculus; Codd Rules.

SQL: Data Definition and Data Types; Constraints, Queries, Insert, Delete, and Update Statements; Views, Stored Procedures and Functions; Database Triggers, SQL Injection.

Normalization for Relational Databases: Functional Dependencies and Normalization; Algorithms for Query Processing and Optimization; Transaction Processing, Concurrency Control Techniques, Database Recovery Techniques, Object and Object-Relational Databases; Database Security and Authorization.

Enhanced Data Models: Temporal Database Concepts, Multimedia Databases, Deductive Databases, XML and Internet Databases; Mobile Databases, Geographic Information Systems, Genome Data Management, Distributed Databases and Client-Server Architectures.

Data Warehousing and Data Mining: Data Modeling for Data Warehouses, Concept Hierarchy, OLAP and OLTP; Association Rules, Classification, Clustering, Regression,

Support Vector Machine, K-Nearest Neighbour, Hidden Markov Model, Summarization, Dependency Modeling, Link Analysis, Sequencing Analysis, Social Network Analysis.

Big Data Systems: Big Data Characteristics, Types of Big Data, Big Data Architecture, Introduction to Map-Reduce and Hadoop; Distributed File System, HDFS.
NOSQL: NOSQL and Query Optimization; Different NOSQL Products, Querying and Managing NOSQL; Indexing and Ordering Data Sets; NOSQL in Cloud.

Unit – 5 : System Software and Operating System

System Software: Machine, Assembly and High-Level Languages; Compilers and Interpreters; Loading, Linking and Relocation; Macros, Debuggers.

Basics of Operating Systems: Operating System Structure, Operations and Services; System Calls, Operating-System Design and Implementation; System Boot.
Process Management: Process Scheduling and Operations; Interprocess Communication, Communication in Client–Server Systems, Process Synchronization, Critical-Section Problem, Peterson's Solution, Semaphores, Synchronization.

Threads: Multicore Programming, Multithreading Models, Thread Libraries, Implicit Threading, Threading Issues.

CPU Scheduling: Scheduling Criteria and Algorithms; Thread Scheduling, Multiple-Processor Scheduling, Real-Time CPU Scheduling.

Deadlocks: Deadlock Characterization, Methods for Handling Deadlocks, Deadlock Prevention, Avoidance and Detection; Recovery from Deadlock.

Memory Management: Contiguous Memory Allocation, Swapping, Paging, Segmentation, Demand Paging, Page Replacement, Allocation of Frames, Thrashing, Memory-Mapped Files.

Storage Management: Mass-Storage Structure, Disk Structure, Scheduling and Management, RAID Structure.

File and Input/Output Systems: Access Methods, Directory and Disk Structure; File-System Mounting, File Sharing, File-System Structure and Implementation; Directory Implementation, Allocation Methods, Free-Space Management, Efficiency and Performance; Recovery, I/O Hardware, Application I/O Interface, Kernel I/O Subsystem, Transforming I/O Requests to Hardware Operations.

Security: Protection, Access Matrix, Access Control, Revocation of Access Rights, Program Threats, System and Network Threats; Cryptography as a Security Tool, User Authentication, Implementing Security Defenses.

Virtual Machines: Types of Virtual Machines and Implementations; Virtualization.

Linux Operating Systems: Design Principles, Kernel Modules, Process Management, Scheduling, Memory Management, File Systems, Input and Output; Interprocess Communication, Network Structure.

Click Here to Buy book on Amazon

Windows Operating Systems: Design Principles, System Components, Terminal Services and Fast User Switching; File System, Networking.

Distributed Systems: Types of Network based Operating Systems, Network Structure, Communication Structure and Protocols; Robustness, Design Issues, Distributed File Systems.

Unit – 6 : Software Engineering

Software Process Models: Software Process, Generic Process Model – Framework Activity, Task Set and Process Patterns; Process Lifecycle, Prescriptive Process Models, Project Management, Component Based Development, Aspect-Oriented Software Development, Formal Methods, Agile Process Models – Extreme Programming (XP), Adptive Software Development, Scrum, Dynamic System Development Model, Feature Driven Development, Crystal, Web Engineering**.**

Software Requirements: Functional and Non-Functional Requirements; Eliciting Requirements, Developing Use Cases, Requirement Analysis and Modelling; Requirements Review, Software Requirment and Specification (SRS) Document.

Software Design: Abstraction, Architecture, Patterns, Separation of Concerns, Modularity, Information Hiding, Functional Independence, Cohesion and Coupling; Object-Oriented Design, Data Design, Architectural Design, User Interface Design, Component Level Design.

Software Quality: McCall's Quality Factors, ISO 9126 Quality Factors, Quality Control, Quality Assurance, Risk Management, Risk Mitigation, Monitoring and Management (RMMM); Software Reliability.

Estimation and Scheduling of Software Projects: Software Sizing, LOC and FP based Estimations; Estimating Cost and Effort; Estimation Models, Constructive Cost Model (COCOMO), Project Scheduling and Staffing; Time-line Charts.

Software Testing: Verification and Validation; Error, Fault, Bug and Failure; Unit and Integration Tesing; White-box and Black-box Testing; Basis Path Testing, Control Structure Testing, Deriving Test Cases, Alpha and Beta Testing; Regression Testing, Performance Testing, Stress Testing.

Software Configuration Management: Change Control and Version Control; Software Reuse, Software Re-engineering, Reverse Engineering.

Unit – 7 : Data Structures and Algorithms

Data Structures: Arrays and their Applications; Sparse Matrix, Stacks, Queues, Priority Queues, Linked Lists, Trees, Forest, Binary Tree, Threaded Binary Tree, Binary Search Tree, AVL Tree, B Tree, B+ Tree, B* Tree, Data Structure for Sets, Graphs, Sorting and Searching Algorithms; Hashing.

Downloaded from www.pdf.tube

Performance Analysis of Algorithms and Recurrences: Time and Space Complexities; Asymptotic Notation, Recurrence Relations.

Design Techniques: Divide and Conquer; Dynamic Programming, Greedy Algorithms, Backtracking, Branch and Bound.

Lower Bound Theory: Comparison Trees, Lower Bounds through Reductions.

Graph Algorithms: Breadth-First Search, Depth-First Search, Shortest Paths, Maximum Flow, Minimum Spanning Trees.

Complexity Theory: P and NP Class Problems; NP-completeness and Reducibility.

Selected Topics: Number Theoretic Algorithms, Polynomial Arithmetic, Fast Fourier Transform, String Matching Algorithms.

Advanced Algorithms: Parallel Algorithms for Sorting, Searching and Merging, Approximation Algorithms, Randomized Algorithms.

Unit – 8 : Theory of Computation and Compilers

Theory of Computation: Formal Language, Non-Computational Problems, Diagonal Argument, Russels's Paradox.

Regular Language Models: Deterministic Finite Automaton (DFA), Non-Deterministic Finite Automaton (NDFA), Equivalence of DFA and NDFA, Regular Languages, Regular Grammars, Regular Expressions, Properties of Regular Language, Pumping Lemma, Non-Regular Languages, Lexical Analysis.

Context Free Language: Pushdown Automaton (PDA), Non-Deterministic Pushdown Automaton (NPDA), Context Free Grammar, Chomsky Normal Form, Greibach Normal Form, Ambiguity, Parse Tree Representation of Derivation Trees, Equivalence of PDA's and Context Free Grammars; Properties of Context Free Language.

Turing Machines (TM): Standard Turing Machine and its Variations; Universal Turing Machines, Models of Computation and Church-Turing Thesis; Recursive and Recursively-Enumerable Languages; Context-Sensitive Languages, Unrestricted Grammars, Chomsky Hierarchy of Languages, Construction of TM for Simple Problems.

Unsolvable Problems and Computational Complexity: Unsolvable Problem, Halting Problem, Post Correspondence Problem, Unsolvable Problems for Context-Free Languages, Measuring and Classifying Complexity, Tractable and Intractable Problems.

Syntax Analysis: Associativity, Precedence, Grammar Transformations, Top Down Parsing, Recursive Descent Predictive Parsing, LL(1) Parsing, Bottom up Parsing, LR Parser, LALR(1) Parser.

Semantic Analysis: Attribute Grammar, Syntax Directed Definitions, Inherited and Synthesized Attributes; Dependency Graph, Evaluation Order, S-attributed and L-attributed Definitions; Type-Checking.

NODIA

Run Time System: Storage Organization, Activation Tree, Activation Record, Stack Allocation of Activation Records, Parameter Passing Mechanisms, Symbol Table.

Intermediate Code Generation: Intermediate Representations, Translation of Declarations, Assignments, Control Flow, Boolean Expressions and Procedure Calls.

Code Generation and Code Optimization: Control-flow, Data-flow Analysis, Local Optimization, Global Optimization, Loop Optimization, Peep-Hole Optimization, Instruction Scheduling.

Unit – 9 : Data Communication and Computer Networks

Data Communication: Components of a Data Communication System, Simplex, Half-Duplex and Duplex Modes of Communication; Analog and Digital Signals; Noiseless and Noisy Channels; Bandwidth, Throughput and Latency; Digital and Analog Transmission; Data Encoding and Modulation Techniques; Broadband and Baseband Transmission; Multiplexing, Transmission Media, Transmission Errors, Error Handling Mechanisms.

Computer Networks: Network Topologies, Local Area Networks, Metropolitan Area Networks, Wide Area Network, Wireless Networks, Internet.

Network Models: Layered Architecture, OSI Reference Model and its Protocols; TCP/IP Protocol Suite, Physical, Logical, Port and Specific Addresses; Switching Techniques.

Functions of OSI and TCP/IP Layers: Framing, Error Detection and Correction; Flow and Error Control; Sliding Window Protocol, HDLC, Multiple Access – CSMA/CD, CSMA/CA, Reservation, Polling, Token Passing, FDMA, CDMA, TDMA, Network Devices, Backbone Networks, Virtual LANs.
IPv4 Structure and Address Space; Classful and Classless Addressing; Datagram, Fragmentation and Checksum; IPv6 Packet Format, Mapping Logical to Physical Address (ARP), Direct and Indirect Network Layer Delivery; Routing Algorithms, TCP, UDP and SCTP Protocols; Flow Control, Error Control and Congestion Control in TCP and SCTP.

World Wide Web (WWW): Uniform Resource Locator (URL), Domain Name Service (DNS), Resolution - Mapping Names to Addresses and Addresses to Names; Electronic Mail Architecture, SMTP, POP and IMAP; TELNET and FTP.

Network Security: Malwares, Cryptography and Steganography; Secret-Key Algorithms, Public-Key Algorithms, Digital Signature, Virtual Private Networks, Firewalls.

Mobile Technology: GSM and CDMA; Services and Architecture of GSM and Mobile Computing; Middleware and Gateway for Mobile Computing; Mobile IP and Mobile Communication Protocol; Communication Satellites, Wireless Networks and Topologies; Cellular Topology, Mobile Adhoc Networks, Wireless Transmission and Wireless LANs; Wireless Geolocation Systems, GPRS and SMS.

Cloud Computing and IoT: SaaS, PaaS, IaaS, Public and Private Cloud; Virtualization, Virtual Server, Cloud Storage, Database Storage, Resource Management, Service Level Agreement, Basics of IoT.

Click Here to Buy book on Amazon

Unit – 10 : Artificial Intelligence (AI)

Approaches to AI: Turing Test and Rational Agent Approaches; State Space Representation of Problems, Heuristic Search Techniques, Game Playing, Min-Max Search, Alpha Beta Cutoff Procedures.

Knowledge Representation: Logic, Semantic Networks, Frames, Rules, Scripts, Conceptual Dependency and Ontologies; Expert Systems, Handling Uncertainty in Knowledge.

Planning: Components of a Planning System, Linear and Non Linear Planning; Goal Stack Planning, Hierarchical Planning, STRIPS, Partial Order Planning.

Natural Language Processing: Grammar and Language; Parsing Techniques, Semantic Analysis and Prgamatics.

Multi Agent Systems: Agents and Objects; Agents and Expert Systems; Generic Structure of Multiagent System, Semantic Web, Agent Communication, Knowledge Sharing using Ontologies, Agent Development Tools.

NODIA

Fuzzy Sets: Notion of Fuzziness, Membership Functions, Fuzzification and Defuzzification; Operations on Fuzzy Sets, Fuzzy Functions and Linguistic Variables; Fuzzy Relations, Fuzzy Rules and Fuzzy Inference; Fuzzy Control System and Fuzzy Rule Based Systems.

Genetic Algorithms (GA): Encoding Strategies, Genetic Operators, Fitness Functions and GA Cycle; Problem Solving using GA.

Artificial Neural Networks (ANN): Supervised, Unsupervised and Reinforcement Learning; Single Perceptron, Multi Layer Perceptron, Self Organizing Maps, Hopfield Network.

Downloaded from www.pdf.tube

CHAPTER 01

Discrete Structures and Optimization

1. Which of the following is a declarative sentence that is either true or false, but not both?
(A) A command or an order
(B) A proposition
(C) A question
(D) An exclamation

2. Consider the following statements: Statement I: The union of two sets A and B, denoted $A \cup B$, is the set of all elements that are in A, or in B, or in both. Statement II: The intersection of two sets A and B, denoted $A \cap B$, is the set of all elements that are in A and not in B.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

3. What does the Pigeonhole Principle state?
(A) If n pigeons are put into m pigeonholes with $n < m$, then at least one pigeonhole must contain more than one pigeon.
(B) The total number of possible permutations of a set containing n distinct objects is exactly n!.
(C) If n pigeons are put into m pigeonholes with $n > m$, then at least one pigeonhole must contain more than one pigeon.
(D) Every complex mathematical problem can be solved by dividing it into smaller, identical subproblems.

4. An algebraic structure (G, *) is called a semi-group if it satisfies which of the following properties?
(A) Associativity and existence of an identity element
(B) Closure, Associativity, and existence of an inverse
(C) Commutativity and Associativity
(D) Closure and Associativity

5. In graph theory, what is a path that starts and ends at the same vertex?
(A) A trail
(B) A walk
(C) A circuit
(D) A simple path

6. The simplification of the Boolean expression A + A'B results in which of the following?
(A) A' + B
(B) A + B
(C) B
(D) A

7. Which of the following is NOT a requirement for a standard linear programming problem model?
(A) An objective function to be maximized or minimized.
(B) A set of linear constraints.
(C) All decision variables must be non-negative.
(D) The objective function must be a quadratic equation.

NODIA

8. Consider the following statements regarding propositional logic. Which statement is incorrect?
(A) The expression $p \rightarrow q$ is a tautology.
(B) The expression $p \vee \neg p$ is a tautology.
(C) The expression $p \wedge \neg p$ is a contradiction.
(D) The expressions $p \rightarrow q$ and $\neg p \vee q$ are logically equivalent.

9. Which property must a relation R on a set A have to be an equivalence relation?
(A) Reflexive and Symmetric only
(B) Symmetric and Transitive only
(C) Reflexive, Symmetric, and Transitive
(D) Reflexive and Anti-symmetric

10. How many different ways can the letters of the word 'LOGIC' be arranged?
(A) 120
(B) 24
(C) 720
(D) 60

11. A group (G, *) is called an Abelian group if it satisfies which additional property?
(A) Associativity
(B) Existence of an identity
(C) Existence of an inverse
(D) Commutativity

12. A graph with no loops and no parallel edges between the same pair of vertices is called a:
(A) Multigraph
(B) Simple graph
(C) Pseudograph
(D) Weighted graph

13. Consider the following statements: Statement I: A Boolean function can be represented by a truth table. Statement II: Karnaugh maps are used for simplifying Boolean functions.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

14. The graphical method is used to solve Linear Programming Problems involving how many decision variables?
(A) One variable
(B) Three variables
(C) Any number of variables
(D) Two variables

15. What is the contrapositive of the implication If it is raining, then the ground is wet?
(A) If the ground is not wet, then it is not raining.
(B) If it is not raining, then the ground is not wet.
(C) If the ground is wet, then it is raining at that time.
(D) It is raining and the ground is not wet currently.

16. Let R be a relation on the set of integers Z defined by aRb if a - b is an even number. Which of the following properties does R possess?
(A) Reflexive and Symmetric, but not Transitive
(B) An Equivalence Relation
(C) Symmetric and Transitive, but not Reflexive
(D) Only Symmetric

17. In a class of 30 students, 15 students play cricket and 20 students play football. If 5 students play both, how many students play at least one of the two sports?
(A) 25
(B) 30
(C) 35
(D) 10

18. In group theory, a one-to-one and onto homomorphism is known as:
(A) An isomorphism
(B) An endomorphism
(C) An automorphism
(D) A semi-group

19. A path in a graph that visits every edge exactly once is known as:
(A) A Hamiltonian path
(B) A spanning path
(C) An Eulerian path
(D) A shortest path

20. A Boolean function that evaluates to 1 for every combination of input values is called a:
(A) Contradiction
(B) Minterm
(C) Maxterm
(D) Tautology

21. In the Simplex method, what is the role of slack variables?
(A) To convert 'greater than or equal to' inequalities into equalities.
(B) To represent the objective function value.
(C) To convert 'less than or equal to' inequalities into equalities.
(D) To ensure all variables are integers.

22. The statement ∀x P(x) is read as:
(A) There exists an x for which P(x) is true.
(B) For all x, P(x) is true.
(C) For some x, P(x) is false.
(D) There exists a unique x for which P(x) is true.

23. Which of the following set operations is analogous to logical OR?
(A) Intersection
(B) Complement
(C) Set Difference

Click Here to Buy book on Amazon

(D) Union

24. What is the first step in a proof by mathematical induction?
(A) The basis step
(B) The inductive step
(C) The assumption step
(D) The conclusion step

25. An algebraic structure (R, +, *) with two binary operations is a Ring if it satisfies all of the following except:
(A) (R, +) is an abelian group.
(B) (R, *) is a semi-group.
(C) Every non-zero element has a multiplicative inverse.
(D) Multiplication is distributive over addition.

26. The number of vertices of odd degree in any graph is always:
(A) Even
(B) Odd
(C) A prime number
(D) Zero

27. The canonical form of a Boolean expression consisting of a sum of minterms is known as:
(A) Sum of Products (SOP)
(B) Product of Sums (POS)
(C) Conjunctive Normal Form (CNF)
(D) Disjunctive Normal Form (DNF)

28. The PERT technique of network analysis is primarily concerned with what aspect of a project?
(A) Cost
(B) Time
(C) Quality
(D) Resources

29. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The proposition $p \wedge (q \vee r)$ is equivalent to $(p \wedge q) \vee (p \wedge r)$.
Reason (R) : This is an example of the distributive law in propositional logic.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

30. A relation R on a set A is a partially ordering if it is:
(A) Reflexive, Symmetric, and Transitive
(B) Symmetric and Transitive
(C) Reflexive and Symmetric
(D) Reflexive, Anti-symmetric, and Transitive

31. If $P(A) = 0.4$, $P(B) = 0.5$, and $P(A \cap B) = 0.2$, what is $P(A \cup B)$?
(A) 0.7
(B) 0.9
(C) 0.6
(D) 0.1

32. A non-empty subset H of a group G is a subgroup of G if and only if for any a, b in H:
(A) $a * b$ is in H
(B) $a * b^{-1}$ is in H
(C) $a^{-1} * b$ is in H
(D) $a * b = b * a$

33. Which algorithm is commonly used to find the shortest path in a weighted graph from a single source vertex to all other vertices?
(A) Prim's algorithm
(B) Dijkstra's algorithm
(C) Kruskal's algorithm
(D) Floyd-Warshall algorithm

34. The dual of the Boolean expression $(X + Y)(X' + Z)$ is:
(A) $XY + X'Z$
(B) $(X'Y') + (X+Z')$
(C) $(X' + Y')(X + Z')$
(D) $X'Y' + XZ'$

35. In the context of transportation models, what is the primary objective?
(A) To maximize the number of operational routes used in the network.
(B) To maximize the profit from a single source by adjusting distributions.
(C) To minimize the total cost of shipping goods from sources to destinations.
(D) To ensure all destinations receive goods from all available sources in the set.

NODIA

36. Which rule of inference is represented by the argument: If P, then Q. P. Therefore, Q?
(A) Modus Ponens
(B) Hypothetical Syllogism
(C) Disjunctive Syllogism
(D) Modus Ponens

37. If A = {1, 2, 3} and B = {a, b}, what is the cardinality of the power set of $A \times B$?
(A) 6
(B) 32
(C) 64
(D) 12

38. The principle of inclusion-exclusion is used for:
(A) Proving statements for all natural numbers.
(B) Finding the number of elements in the union of two or more sets.
(C) Arranging objects in a sequence.
(D) Selecting a subset of objects from a set.

39. An integral domain is a commutative ring with a unit element that also satisfies which property?
(A) Every element has a multiplicative inverse.
(B) It is a finite set.
(C) The two operations are identical.
(D) It has no zero divisors.

40. A graph is called bipartite if and only if:
(A) It has no cycles of odd length.
(B) It has an Eulerian circuit.
(C) It is planar.
(D) It is a complete graph.

41. Which of the following is not a valid representation of a Boolean function?
(A) Truth Table
(B) Karnaugh Map
(C) Venn Diagram
(D) PERT Chart

42. What does sensitivity analysis in Linear Programming study?
(A) The effect of changes in input parameters on the optimal solution.
(B) The sensitivity of the computer to the complexity of the problem size.
(C) The systematic selection of the most sensitive variables in the model.
(D) The mathematical process of converting the primal problem to its dual.

NODIA

43. Which of the following statements about logical connectives is correct?
(A) $p \leftrightarrow q$ is true only when p and q have different truth values.
(B) $p \wedge q$ is true when either p or q is true.
(C) $p \vee q$ is false when either p or q is true.
(D) $p \rightarrow q$ is false only when p is true and q is false.

44. A relation where every element is related to itself is called:
(A) Symmetric
(B) Reflexive
(C) Transitive
(D) Anti-symmetric

45. Bayes' theorem is used to calculate:
(A) The total number of unique permutations of a set.
(B) The mathematical sum of an infinite geometric series.
(C) The probability of an event based on prior knowledge of conditions.
(D) The weight of the shortest path in a directed graph.

46. What is a homomorphism from a group to itself called?
(A) Isomorphism
(B) Automorphism
(C) Monomorphism
(D) Endomorphism

47. A spanning tree of a connected graph G is a subgraph that:
(A) Contains all the vertices of G and is a tree.
(B) Is a complete graph.
(C) Contains all the edges of G.
(D) Contains exactly one cycle.

48. A minterm is a product term that contains every variable of a function exactly once, either in its complemented or uncomplemented form. What is the minterm corresponding to the binary input 101 for variables A, B, C?
(A) A'BC
(B) AB'C
(C) A'B'C
(D) ABC'

49. An optimal solution to an assignment problem can be found using which method?
(A) Simplex Method
(B) Hungarian Method
(C) Graphical Method
(D) Vogel's Approximation Method

50. Match the items in List I with those in List II:
List I (Logical Laws):
P. $p \vee \neg p \equiv T$
Q. $p \wedge q \equiv q \wedge p$
R. $\neg(p \wedge q) \equiv \neg p \vee \neg q$
S. $p \vee (p \wedge q) \equiv p$.
List II (Name of Law):
1. De Morgan's Law
2. Commutative Law
3. Absorption Law
4. Complement Law.
(A) P-4, Q-2, R-1, S-3
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-4, R-1, S-3

51. What is a set with a partially ordering relation defined on it called?
(A) A group
(B) A ring
(C) A poset
(D) A lattice

52. Consider the following statements: Statement I: Permutation refers to the selection of objects where order does not matter. Statement II: Combination refers to the arrangement of objects where order matters.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

53. An isomorphism from a group G to itself is called:
(A) An endomorphism
(B) A homomorphism
(C) An automorphism
(D) A subgroup

54. The problem of finding a circuit in a graph that visits every vertex exactly once is the:
(A) Eulerian circuit problem
(B) Hamiltonian circuit problem
(C) Chinese postman problem
(D) Graph coloring problem

55. Which of the following is an example of a universal logic gate, from which any Boolean function can be implemented?
(A) AND gate
(B) OR gate
(C) XOR gate
(D) NAND gate

56. The Dual Simplex Method is used when:
(A) The initial simplex tableau is optimal but infeasible.
(B) The problem has more than two variables.
(C) The objective function is to be maximized.
(D) The initial solution is feasible but not optimal.

57. The process of deriving new true statements from existing true statements is known as:
(A) Quantification
(B) Predication
(C) Inference
(D) Equivalence

58. Which of the following statements about sets is incorrect?
(A) The order of elements in a set is important.
(B) A set can be an element of another set.
(C) The empty set is a subset of every set.
(D) The cardinality of the power set of a set with n elements is 2n.

59. In the context of probability, two events are said to be mutually exclusive if:
(A) Their union is the sample space.
(B) They are independent.
(C) The probability of one event depends on the other.
(D) Their intersection is the empty set.

60. A ring in which multiplication is commutative is called a:
(A) Division ring
(B) Commutative ring
(C) Integral domain
(D) Field

61. What is the chromatic number of a graph?
(A) The maximum number of colors needed to

Downloaded from www.pdf.tube

color the edges of the graph.
(B) The total number of vertices that are contained in the graph set.
(C) The minimum number of colors needed to color the vertices such that no two adjacent vertices share the same color.
(D) The total number of fundamental cycles that exist in the graph.

62. The representation of a Boolean function as a Product of Sums (POS) corresponds to which normal form?
(A) Disjunctive Normal Form (DNF)
(B) Conjunctive Normal Form (CNF)
(C) Principal Disjunctive Normal Form
(D) Principal Conjunctive Normal Form

63. Which of the following is a key component of Integer Programming?
(A) Some or all of the decision variables are restricted to be integer values.
(B) All variables can take any real value.
(C) The objective function must be non-linear.
(D) The constraints must be non-linear.

64. Consider the following statements:
(i) It must have an identity element.
(ii) It must be commutative.
(iii) Every element must have an inverse.
(iv) It must satisfy the associative law.
Which of the above statement(s) is/are correct requirements for a group?
(A) (i), (ii), (iii)
(B) (i), (iii), (iv)
(C) (i), (ii), (iv)
(D) All of the above

65. A predicate is a property or relationship that can be affirmed or denied about the objects of a certain domain. It becomes a proposition when:
(A) It is negated.
(B) Its variables are bound by quantifiers or assigned specific values.
(C) It is combined with another predicate using a logical connective.
(D) It is written in a normal form.

66. A relation R on set A is anti-symmetric if:
(A) $(a, b) \in R$ and $(b, a) \in R$ implies $a = b$ for all $a, b \in A$.
(B) $(a, b) \in R$ implies $(b, a) \in R$ for all $a, b \in A$.
(C) $(a, a) \notin R$ for all $a \in A$.
(D) $(a, b) \in R$ and $(b, c) \in R$ implies $(a, c) \in R$.

67. What is the value of C(n, 1)?
(A) 1
(B) n-1
(C) n
(D) n!

68. What is the critical path in a PERT/CPM network?
(A) The shortest path available from the start node to the end node.
(B) The path consisting of the maximum number of individual activities.
(C) The path that involves the most critical resources in the model.
(D) The longest path in the network, which determines the minimum project completion time.

69. A graph that can be drawn in a plane without any edges crossing is called a:
(A) Complete graph
(B) Bipartite graph
(C) Planar graph
(D) Weighted graph

70. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : In any group, the identity element is its own inverse.
Reason (R) : For an identity element 'e', $e * e = e$, which satisfies the property $a * a^{-1} = e$.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

71. A maxterm is a sum term that contains every variable of a function exactly once, either in its complemented or uncomplemented form. What is the maxterm corresponding to the binary input 010 for variables A, B, C?
(A) A'+B+C'
(B) A'+B'+C
(C) A+B+C'
(D) A+B'+C

NODIA

72. In the inductive step of a proof by mathematical induction for a statement P(n), what is assumed?
(A) P(k) is true for some arbitrary integer k.
(B) P(1) is true.
(C) P(k+1) is true.
(D) P(n) is true for all integers n.

73. The logical equivalence $p \rightarrow q \equiv \neg p \vee q$ is known as:
(A) De Morgan's Law
(B) Commutative Law
(C) Implication Equivalence
(D) Idempotent Law

74. A connected graph G has an Eulerian circuit if and only if every vertex of G has:
(A) Even degree
(B) Odd degree
(C) A degree of at least 2
(D) The same degree

75. What is the primary purpose of tree traversal algorithms like in-order, pre-order, and post-order?
(A) To find the shortest path between two nodes.
(B) To determine if the tree is balanced.
(C) To calculate the height of the tree.
(D) To visit every node in the tree in a systematic way.

76. Which of the following is a key characteristic of a field in abstract algebra?
(A) It is a commutative ring where multiplication is not associative.
(B) It is a commutative ring in which every non-zero element has a multiplicative inverse.
(C) It is a non-commutative ring with a unit element.
(D) It is an integral domain that is infinite.

77. The assignment problem is a special case of which of the following problems?
(A) Integer programming problem
(B) Shortest path problem
(C) Transportation problem
(D) Knapsack problem

78. A proposition that is always false is known as:
(A) A tautology
(B) A contingency
(C) An equivalence
(D) A contradiction

79. The set of all subsets of a set A is called:
(A) The power set of A
(B) The universal set
(C) The null set
(D) The complement of A

80. How many permutations are there of n distinct objects taken r at a time, denoted P(n, r)?
(A) n! / r!
(B) n! / (n-r)!
(C) n! / (r! * (n-r)!)
(D) (n-r)!

81. Consider the following statements about group theory:
(i) A group can have more than one identity element.
(ii) The inverse of an element in a group is unique.
(iii) The set of integers with addition forms a group.
(iv) The set of natural numbers with addition forms a group.
(v) The cancellation laws hold in a group.
Which of the above statement(s) is/are correct?
(A) (i), (iv) only
(B) (ii), (iii), (v) only
(C) (i), (iii), (v) only
(D) (ii), (iv) only

82. In a rooted tree, what is a vertex with no children called?
(A) A leaf
(B) A root
(C) A parent
(D) An internal node

83. The simplification of a Boolean function using a Karnaugh map is based on the principle of combining terms that differ in:
(A) Two variables.
(B) All variables.
(C) Exactly one variable.
(D) No variables.

84. What is the term for the slack on the critical path in a project schedule?
(A) Positive
(B) Negative

Click Here to Buy book on Amazon

(C) Infinite
(D) Zero or negative

85. Which of the following represents the nested quantifier concept?
(A) $\forall x\ P(x)$
(B) $\exists x\ P(x)$
(C) $\forall x\ \exists y\ P(x, y)$
(D) $P(x) \rightarrow Q(x)$

86. Which of the following is NOT a property of an equivalence relation?
(A) Reflexive
(B) Anti-symmetric
(C) Symmetric
(D) Transitive

87. If you have 5 distinct books, how many different combinations of 3 books can you choose?
(A) 20
(B) 60
(C) 15
(D) 10

88. What is a cut-set in a graph?
(A) A set of edges whose removal increases the number of connected components.
(B) A set of vertices whose removal disconnects the graph.
(C) A path that visits every vertex exactly once.
(D) A cycle that includes all edges.

89. A code where no codeword is a prefix of another codeword is known as a:
(A) Binary code
(B) Error-correcting code
(C) Prefix code
(D) Gray code

90. In a linear programming problem, the feasible region is the set of points satisfying:
(A) All constraints, including the non-negativity restrictions.
(B) Only the objective function.
(C) Only the non-negativity restrictions.
(D) Only the main constraints.

91. Match the items in List I with those in List II:
List I (Algebraic Structure): P. Semi-group, Q. Monoid, R. Group, S. Abelian Group.
List II (Properties): 1. Closure, Associativity, Identity, Inverse; 2. Closure, Associativity; 3. Closure, Associativity, Identity; 4. Closure, Associativity, Identity, Inverse, Commutativity.
(A) P-2, Q-3, R-1, S-4
(B) P-3, Q-2, R-4, S-1
(C) P-2, Q-1, R-3, S-4
(D) P-1, Q-2, R-3, S-4

NODIA

92. The logical connective for conjunction is represented by which symbol?
(A) $\vee$
(B) $\wedge$
(C) $\rightarrow$
(D) $\neg$

93. Consider a relation R on the set of positive integers where aRb if 'a divides b'. Which properties does this relation have?
(A) Reflexive, Symmetric, Transitive
(B) Symmetric, Transitive
(C) Reflexive, Anti-symmetric, Transitive
(D) Reflexive, Symmetric

94. What is the probability of drawing an ace from a standard deck of 52 playing cards?
(A) 1/52
(B) 4/13
(C) 2/13
(D) 1/13

95. In a weighted graph, the shortest path problem aims to find a path between two vertices such that:
(A) The sum of the weights of its constituent edges is minimized.
(B) The sum of the weights of its constituent edges is maximized.
(C) The number of edges is minimized.
(D) The path passes through all vertices.

96. A Boolean function can be uniquely represented by its:
(A) Simplified SOP form
(B) Canonical Sum of Products (Disjunctive Normal Form)
(C) Number of literals
(D) K-map

97. What is resource levelling in project management?

(A) Increasing the project budget.
(B) Adjusting the schedule to reduce peaks and troughs in resource demand.
(C) Reducing the project duration.
(D) Identifying the critical path.

98. A normal form of a propositional formula is an equivalent formula that:
(A) Has a specific syntactic structure, like CNF or DNF.
(B) Uses only AND and NOT operators.
(C) Is always a tautology.
(D) Has the minimum number of variables.

99. Which of the following is NOT a valid set operation?
(A) Union
(B) Intersection
(C) Division
(D) Complement

100. In a group (G, *), if a * b = e (the identity element), then 'b' is the what of 'a'?
(A) Subgroup
(B) Homomorphism
(C) Identity
(D) Inverse

101. A complete graph with 'n' vertices, Kn, has how many edges?
(A) n
(B) n(n-1)
(C) n(n-1)/2
(D) n-1

102. Consider the following statements: Statement I: Post-order traversal of a binary search tree visits the root node last. Statement II: Pre-order traversal of a binary search tree visits the root node first.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

103. The term 'resource levelling' in PERT-CPM refers to the process of:
(A) Minimizing project duration by allocating more resources.
(B) Calculating the critical path of the project.
(C) Reducing the total cost of the project by using cheaper resources.
(D) Ensuring that resource usage is kept at a constant level during the project.

104. What is the negation of the statement For all integers x, $x > 5$?
(A) There exists an integer x such that $x \leq 5$.
(B) For all integers x, $x \leq 5$.
(C) There exists an integer x such that $x > 5$.
(D) For all integers x, $x < 5$.

105. An equivalence relation on a set partitions the set into:
(A) A partially ordered set
(B) A single subset
(C) Disjoint equivalence classes
(D) A totally ordered set

106. In how many ways can a committee of 3 be formed from a group of 7 people?
(A) 35
(B) 21
(C) 210
(D) 42

107. Which of the following is a correct statement about subgroups?
(A) The intersection of any two subgroups of a group G is not necessarily a subgroup of G.
(B) The union of any two subgroups of a group G is always a subgroup of G.
(C) A subgroup must be an abelian group.
(D) Every group has at least two subgroups: the group itself and the trivial subgroup.

108. Which traversal of a tree would result in a sorted list of nodes if the tree is a Binary Search Tree?
(A) Pre-order
(B) In-order
(C) Post-order
(D) Level-order

109. The number of 1s in the binary representation of a number is referred to as its:
(A) Hamming distance
(B) Parity
(C) Hamming weight
(D) Checksum

Downloaded from www.pdf.tube

110. Which of the following conditions must be met for a graph to have an Eulerian path but not an Eulerian circuit?
(A) All vertices have even degree.
(B) All vertices have odd degree.
(C) The graph is disconnected.
(D) Exactly two vertices have odd degree.

111. Consider the following statements about Boolean algebra:
(i) $A + 1 = 1$, (ii) $A . 0 = 0$, (iii) $A + A' = 1$, (iv) $A . A' = 0$, (v) $A + A = A$.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iv) only
(B) (i), (iii), (v) only
(C) All (i), (ii), (iii), (iv), and (v) are correct.
(D) (ii), (iii), (iv) only

112. The mathematical model of a linear programming problem consists of which of the following components?
(A) Objective function, Slack variables, Artificial variables
(B) Decision variables, Objective function, Constraints
(C) Decision variables, Feasible region, Optimal solution
(D) Constraints, Non-negativity restrictions, Corner points

113. Arrange the following steps of the Simplex method in the correct chronological order for a maximization problem:
1. Construct the initial simplex tableau.
2. Select the pivot element.
3. Check for optimality. If optimal, stop.
4. Identify the entering and leaving variables.
5. Perform row operations to make the pivot element 1 and other elements in the pivot column 0.

(A) 1, 3, 4, 2, 5
(B) 1, 4, 2, 5, 3
(C) 1, 2, 3, 4, 5
(D) 3, 1, 4, 2, 5

114. A compound proposition that is neither a tautology nor a contradiction is called a:
(A) Contingency
(B) Fallacy
(C) Predicate
(D) Quantifier

NODIA

115. Let R be a relation on the set of people such that $(a,b) \in R$ if 'a is a sister of b'. Which statement best describes the relation R?
(A) R is symmetric.
(B) R is an equivalence relation.
(C) R is not symmetric but is transitive.
(D) R is symmetric only if we consider a set of all females.

116. The basis step and the inductive step are the two main parts of a proof technique known as:
(A) Proof by contradiction
(B) Proof by contraposition
(C) Direct proof
(D) Mathematical induction

117. A subgroup H of a group G is a normal subgroup if for all h in H and all g in G:
(A) $gh = hg$
(B) $gh \in H$
(C) $ghg^{-1} \in H$
(D) $g^{-1}h \in H$

118. A tree is a connected graph that:
(A) has more edges than vertices
(B) has no cycles
(C) has exactly one cycle
(D) is not planar

119. A Boolean function can be simplified using which of the following laws?
(A) Associative and Commutative laws
(B) Distributive law
(C) Idempotent and Absorption laws
(D) All of the above

120. In PERT, the expected time (te) for an activity is calculated using the formula: (where to = optimistic, tm = most likely, tp = pessimistic)
(A) (to + 4tm + tp) / 6
(B) (to + tm + tp) / 3
(C) (to + 2tm + tp) / 4
(D) (to + 6tm + tp) / 8

The rule of inference $\neg q, p \rightarrow q \therefore \neg p$ is known as:
(A) Modus Ponens
(B) Resolution
(C) Modus Tollens
(D) Hypothetical Syllogism

121. If A and B are two sets, the set difference A - B is defined as:
(A) $\{x \mid x \in A \text{ and } x \notin B\}$
(B) $\{x \mid x \in A \text{ and } x \in B\}$
(C) $\{x \mid x \in A \text{ or } x \in B\}$
(D) $\{x \mid x \in B \text{ and } x \notin A\}$

122. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : If k pigeons are placed in n holes and $k > n$, then at least one hole contains more than one pigeon.
Reason (R) : This is the generalized form of the Pigeonhole principle.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

123. Which of the following structures is a field?
(A) The set of integers under addition and multiplication.
(B) The set of rational numbers under addition and multiplication.
(C) The set of natural numbers under addition and multiplication.
(D) The set of 2x2 matrices under matrix addition and multiplication.

124. The problem of finding a minimum weight spanning tree in a weighted, connected, undirected graph can be solved by which of the following algorithms?
(A) Dijkstra's algorithm
(B) Floyd-Warshall algorithm
(C) Kruskal's algorithm
(D) Bellman-Ford algorithm

125. A function f(A, B, C) = A'BC + AB'C + ABC' + ABC is an example of a Boolean function in which form?
(A) Product of Sums (POS) form
(B) Conjunctive Normal Form (CNF)
(C) A simplified form
(D) Sum of Products (SOP) form

126. The difference between PERT and CPM is that:
(A) PERT uses a probabilistic time estimate while CPM uses a deterministic time estimate.
(B) PERT is activity-oriented while CPM is event-oriented.
(C) PERT is used for research projects while CPM is used for construction projects.
(D) PERT considers cost while CPM does not.

127. A proposition $p \leftrightarrow q$ is logically equivalent to:
(A) $(p \rightarrow q) \vee (q \rightarrow p)$
(B) $(p \rightarrow q) \wedge (q \rightarrow p)$
(C) $(p \wedge q) \rightarrow (p \vee q)$
(D) $(p \vee q) \rightarrow (p \wedge q)$

128. Let A = {a, b, c}. Which of the following is NOT a valid relation on A?
(A) {(a, a), (b, b), (c, c)}
(B) {(a, b), (c, d)}
(C) {(a, b), (b, c), (c, a)}
(D) {} (the empty set)

129. What is the fundamental assumption of the inductive hypothesis in a proof by mathematical induction?
(A) The statement P(k) is true for an arbitrary but fixed integer k.
(B) The statement P(n) is true for all $n \geq 1$.
(C) The statement P(1) is true.
(D) The statement P(k+1) must be proven to be true.

130. Match the items in List I with those in List II:
List I (Graph Traversal/Algorithm):
P. Breadth-First Search (BFS)
Q. Depth-First Search (DFS)
R. Dijkstra's Algorithm
S. Prim's Algorithm.
List II (Data Structure Used):
1. Stack
2. Queue
3. Priority Queue (for minimum spanning tree)
4. Priority Queue (for shortest path).
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-1, R-3, S-4
(D) P-1, Q-2, R-4, S-3

131. What is an integral domain?
(A) A ring where every non-zero element has a multiplicative inverse.

Click Here to Buy book on Amazon

(B) A non-commutative ring with unity.
(C) A group that is also a ring.
(D) A commutative ring with unity and no zero-divisors.

132. Consider the following statements:
Statement I : A simple graph can have multiple edges between the same two vertices.
Statement II : A multigraph can have loops.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

133. Cost consideration in project scheduling primarily involves analyzing the trade-off between:
(A) Time and resources
(B) Time and cost
(C) Quality and scope
(D) Risk and reward

134. For predicates M(x): x is a mammal and A(x): x is an animal, the statement All mammals are animals can be written in predicate logic as:
(A) $\forall x\,(M(x) \wedge A(x))$
(B) $\exists x\,(M(x) \rightarrow A(x))$
(C) $\exists x\,(M(x) \wedge A(x))$
(D) $\forall x\,(M(x) \rightarrow A(x))$

135. The relation is a subset of on the power set of a set is an example of:
(A) A partial ordering
(B) An equivalence relation
(C) A total ordering
(D) A symmetric relation

136. If an event can occur in 'm' ways and after it has occurred, another event can occur in 'n' ways, then the total number of ways of occurrence of the events in the given order is:
(A) m + n
(B) m - n
(C) m * n
(D) m / n

137. What does a quotient of an algebraic structure typically refer to?
(A) An algebraic structure formed by partitioning a larger structure by an equivalence relation.
(B) The mathematical process of dividing one element by another in the same structure.
(C) A specific substructure that is also identified as a group in the hierarchy.
(D) A one-to-one homomorphism from a given algebraic structure to itself.

138. Which of the following is true for a bipartite graph?
(A) The graph must be complete and contain no cycles.
(B) The graph must contain at least one triangle of edges.
(C) Every vertex in the graph must have the exact same degree.
(D) The vertices can be partitioned into two disjoint sets, V1 and V2, such that every edge connects a vertex in V1 to one in V2.

139. Which of the following statements about Boolean simplification is incorrect?
(A) $x + xy = x$
(B) $(x + y)' = x' + y'$
(C) $x(x + y) = x$
(D) $x + x'y = x + y$

140. The Hungarian method is an algorithm used to solve:
(A) Transportation problems
(B) Shortest path problems
(C) Assignment problems
(D) Integer programming problems

141. Consider the following statements:
(i) A simple graph with n vertices can have at most n(n-1)/2 edges.
(ii) A complete graph is always planar.
(iii) The sum of degrees of all vertices in a graph is equal to twice the number of edges.
(iv) A tree with n vertices has n edges.
(v) Every connected graph has a spanning tree.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iii)
(B) (i), (iii), (v)
(C) (ii), (iv), (v)
(D) (i), (iv), (v)

142. Which of the following represents De Morgan's Law for quantifiers?
(A) $\neg\forall x\, P(x) \equiv \exists x\, \neg P(x)$
(B) $\neg\exists x\, P(x) \equiv \forall x\, P(x)$

Downloaded from www.pdf.tube

(C) $\forall x\ P(x) \equiv \neg\exists x\ \neg P(x)$
(D) $\exists x\ P(x) \equiv \neg\forall x\ P(x)$

143. The transitive closure of a relation R is the:
(A) Largest relation contained in R that is transitive.
(B) Smallest relation containing R that is transitive.
(C) Smallest reflexive and symmetric relation containing R.
(D) Smallest equivalence relation containing R.

144. The Inclusion-Exclusion Principle provides a way to compute the:
(A) Cardinality of the intersection of sets.
(B) Cardinality of the union of sets.
(C) Number of permutations of a set.
(D) Probability of an event.

145. A semi-group with an identity element is called a:
(A) Monoid
(B) Group
(C) Subgroup
(D) Ring

146. A tree with 'n' vertices has exactly how many edges?
(A) n
(B) n+1
(C) n-1
(D) n/2

147. The dual of a Boolean expression is obtained by:
(A) Interchanging AND and OR operators, and complementing each variable.
(B) Interchanging variables with their complements.
(C) Complementing the entire expression.
(D) Interchanging 0s and 1s, and interchanging AND and OR operators.

148. In a transportation problem, a solution is feasible if it satisfies:
(A) The objective function only
(B) The rim (supply and demand) conditions only
(C) Both the rim conditions and non-negativity constraints
(D) The optimality test

149. Arrange the following algebraic structures in descending order of the number of axioms they must satisfy (most to least):
1. Group, 2. Field, 3. Ring, 4. Semi-group.
(A) 2, 3, 1, 4
(B) 2, 1, 3, 4
(C) 1, 2, 3, 4
(D) 4, 1, 3, 2

NODIA

150. Which of the following is NOT a rule of inference in propositional logic?
(A) Modus Ponens
(B) Resolution
(C) Conjunction Introduction
(D) Existential Instantiation

151. What is the cause-and-effect relationship between a relation being an equivalence relation and the set being partitioned?
(A) An equivalence relation on a set causes the set to be partitioned into disjoint equivalence classes.
(B) If a set is partitioned, it causes the relation on it to become an equivalence relation.
(C) There is no cause-and-effect relationship existing between these two concepts.
(D) A partition causes the relation to be strictly anti-symmetric in nature.

152. Given a bag with 4 red and 6 blue balls, what is the probability of drawing a red ball?
(A) 4/6
(B) 6/10
(C) 4/10
(D) 1/10

153. What is the identity element for the set of integers under the operation of addition?
(A) 0
(B) 1
(C) -1
(D) There is no identity element.

154. What does the term graph coloring refer to?
(A) Assigning colors to edges such that adjacent edges have different colors.
(B) Drawing a graph using different colored pens.
(C) Finding the most colorful path in a graph.
(D) Assigning colors to vertices such that adjacent vertices have different colors.

155. In a K-map, a group of four adjacent 1s corresponds to a term with:
(A) One variable eliminated
(B) Two variables eliminated
(C) Four variables eliminated
(D) No variables eliminated

156. Which of the following statements about integer programming is true?
(A) It is always easier to solve than a corresponding linear programming problem.
(B) The feasible region is always a convex set.
(C) The optimal solution to the relaxed LP problem provides a bound for the integer solution.
(D) It can be solved efficiently using the graphical method for any number of variables.

157. The statement form $(p \wedge q) \rightarrow (p \vee q)$ is a:
(A) Contradiction
(B) Contingency
(C) Quantifier
(D) Tautology

158. The number of tuples in a relation is called its:
(A) Cardinality
(B) Degree
(C) Domain
(D) Attribute

159. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The number of permutations of n distinct objects is n!.
Reason (R) : This is because the first object can be chosen in n ways, the second in n-1 ways, and so on, which is an application of the product rule of counting.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

160. A non-empty set A is termed as an algebraic structure with respect to a binary operation '*' if:
(A) (A, *) is associative
(B) (A, *) satisfies the closure property
(C) (A, *) has an identity element
(D) (A, *) is commutative

161. A graph where there is a weight or cost associated with each edge is known as a:
(A) Weighted Graph
(B) Multigraph
(C) Simple Graph
(D) Directed Graph

162. The simplification of the Boolean function $F(x, y, z) = \Sigma(0, 1, 2, 3)$ results in:
(A) y'
(B) x'y'
(C) x'
(D) z'

163. What is the main objective of the dual simplex method?
(A) To start with an infeasible but better than optimal solution and move towards feasibility.
(B) To handle integer variables in a complex linear programming problem.
(C) To solve transportation problems more efficiently using the rim conditions.
(D) To start with an infeasible but better than optimal solution and move towards feasibility.

164. Match the items in List I with those in List II:
List I (Optimization Problem) :
P. Assignment Problem
Q. Transportation Problem
R. Integer Programming
S. PERT/CPM.
List II (Primary Objective) :
1. Minimize cost of shipping from sources to destinations
2. Assign tasks to workers at minimum cost
3. Project scheduling and management
4. LP with integer-valued decision variables.
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-4, Q-3, R-2, S-1
(D) P-2, Q-4, R-1, S-3

165. A universally quantified statement $\forall x\ P(x)$ is true if and only if:
(A) P(x) is true for at least one value of x in the domain.

Click Here to Buy book on Amazon

(B) P(x) is true for all values of x in the domain.
(C) P(x) is false for at least one value of x in the domain.
(D) P(x) is false for all values of x in the domain.

166. Which of the following is the best description of a partially ordered set (poset)?
(A) A set where every pair of elements in the set is comparable.
(B) A set with a relation that is reflexive, symmetric, and transitive.
(C) A set where no two elements in the collection are related.
(D) A set with a relation that is reflexive, anti-symmetric, and transitive.

167. If a coin is tossed 3 times, what is the probability of getting exactly 2 heads?
(A) 3/8
(B) 1/8
(C) 2/8
(D) 4/8

168. What is a key application of group theory?
(A) Solving linear equations
(B) Database design
(C) Used in cryptography and coding theory
(D) Project scheduling

169. Consider the following properties:
(i) It is connected.
(ii) It contains no cycles.
(iii) There is a unique path between any two vertices.
(iv) The number of vertices is one more than the number of edges.
(v) It must have a designated root.
Which of the above are properties of a general tree?
(A) (i), (ii), (iii) only
(B) (i), (ii), (iii), (iv)
(C) All (i), (ii), (iii), (iv), (v)
(D) (i), (v) only

170. The representation of a Boolean function F as F = Π M, where M are maxterms, is the:
(A) Canonical Sum of Products
(B) Minimal Sum of Products
(C) Minimal Product of Sums
(D) Canonical Product of Sums

171. Consider a linear programming problem. What is the most likely outcome of adding a redundant constraint?
(A) The feasible region will become smaller.
(B) The feasible region and optimal solution will remain unchanged.
(C) The optimal solution will change.
(D) The problem will become infeasible.

172. The logical expression $(P \rightarrow Q) \wedge (Q \rightarrow R) \rightarrow (P \rightarrow R)$ represents which rule of inference?
(A) Modus Ponens
(B) Disjunctive Syllogism
(C) Hypothetical Syllogism
(D) Simplification

173. In set theory, the symmetric difference of two sets A and B, denoted by $A \Delta B$, is defined as:
(A) $(A - B) \cap (B - A)$
(B) $(A \cup B) \cap (A \cap B)$
(C) $A \cap B'$
(D) $(A \cup B) - (A \cap B)$

174. Which statement is true about mathematical induction?
(A) It is used to prove statements about the set of natural numbers.
(B) It can be used to prove results for real numbers.
(C) It is a form of deductive reasoning.
(D) The inductive hypothesis assumes the statement is true for all n.

175. Which of the following statements about rings, integral domains, and fields is correct?
(A) Every ring is an integral domain.
(B) Every field is an integral domain.
(C) Every integral domain is a field.
(D) Every ring is a field.

176. A prefix code is also known as:
(A) An instantaneous code
(B) A cyclic code
(C) A block code
(D) An error-detecting code

NODIA

177. Consider the following statements:
Statement I : The simplification of Boolean functions aims to reduce the number of literals and terms.
Statement II : Reducing literals and terms

generally leads to a simpler and less expensive circuit implementation.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

178. In an assignment problem, if the number of rows is not equal to the number of columns, the problem is said to be:
(A) Infeasible
(B) Degenerate
(C) Unbalanced
(D) Optimal

179. Which of the following is not a standard normal form in propositional logic?
(A) Disjunctive Normal Form (DNF)
(B) Conjunctive Normal Form (CNF)
(C) Principal Disjunctive Normal Form (PDNF)
(D) Transitive Normal Form (TNF)

180. A relation R is defined on the set of integers as aRb if $|a| = |b|$. Which of the following correctly describes this relation?
(A) It is reflexive and symmetric, but not transitive.
(B) It is symmetric and transitive, but not reflexive.
(C) It is an equivalence relation.
(D) It is only reflexive.

181. According to the Pigeonhole Principle, if you have 13 people in a room, what is the minimum number of people who must share the same birth month?
(A) 1
(B) 2
(C) 3
(D) 12

182. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A Hamiltonian circuit in a graph is a circuit that visits every vertex exactly once.
Reason (R) : Not all connected graphs have a Hamiltonian circuit.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

183. The product of algebraic structures refers to:
(A) Creating a new structure from two or more given structures, typically on their Cartesian product.
(B) Multiplying elements within a single structure.
(C) A specific element obtained by a binary operation.
(D) The process of finding a subgroup.

184. In a rooted tree, the level of a node is defined as:
(A) The number of children it has.
(B) The number of siblings it has.
(C) Its distance from the root.
(D) Its depth plus its height.

185. The Boolean function $F(x, y) = x'y' + x'y + xy'$ can be simplified to:
(A) $x' + y'$
(B) $x + y'$
(C) $x' + y$
(D) $x + y$

186. The purpose of introducing artificial variables in the Simplex method is:
(A) To convert the problem from maximization to minimization.
(B) To ensure all variables are integers.
(C) To handle unbounded solutions.
(D) To find an initial basic feasible solution for constraints of type '$\geq$' or '$=$'.

187. Which of the following is a tautology?
(A) $p \wedge q$
(B) $p \rightarrow (p \vee q)$
(C) $p \oplus q$ (XOR)
(D) $\neg p$

188. Which of the following represents the correct order of precedence for logical operators (highest to lowest)?
(A) $\rightarrow, \leftrightarrow, \neg, \wedge, \vee$
(B) $\wedge, \vee, \neg, \rightarrow, \leftrightarrow$
(C) $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
(D) $\neg, \vee, \wedge, \rightarrow, \leftrightarrow$

Downloaded from www.pdf.tube

189. How many edges are in a tree with 100 vertices?
(A) 100
(B) 101
(C) 4950
(D) 99

190. Arrange the following steps for proving a statement P(n) for $n \geq n_0$ using mathematical induction in chronological order:
1. Show that $P(k) \rightarrow P(k+1)$ is true for all $k \geq n_0$.
2. Assume P(k) is true for an arbitrary integer $k \geq n_0$.
3. Conclude that P(n) is true for all $n \geq n_0$. 4. Show that $P(n_0)$ is true.

(A) 4, 2, 1, 3
(B) 1, 2, 3, 4
(C) 4, 1, 2, 3
(D) 2, 1, 4, 3

191. Which of the following is a key property of a semi-group?
(A) Commutativity
(B) Associativity
(C) Existence of identity
(D) Existence of inverse

192. The purpose of a cut-set in a graph is related to the concept of:
(A) Coloring and partitioning
(B) Connectivity and flow
(C) Traversal and searching
(D) Planarity and embedding

193. What is the primary use of a Karnaugh Map?
(A) To simplify a Boolean function.
(B) To represent a Boolean function.
(C) To implement a Boolean function using logic gates.
(D) To find the dual of a Boolean function.

194. Consider the following statements:
(i) PERT is probabilistic and suitable for R&D projects.
(ii) CPM is deterministic and suitable for construction projects.
(iii) The critical path is the shortest path in the network diagram.
(iv) Slack is the amount of time an activity can be delayed without affecting the project completion time.

Which of the above statement(s) is/are correct?
(A) (i), (ii), (iii)
(B) (i), (ii), (iv)
(C) (i), (iii), (iv)
(D) (ii), (iii), (iv)

195. The statement ∃x∀y P(x, y) implies which of the following?
(A) For all y, there is some x such that P(x,y) is true.
(B) For all x, there is some y such that P(x,y) is true.
(C) There is a y such that for all x, P(x,y) is true.
(D) There is an x such that for all y, P(x,y) is true.

196. A relation R on a set A is transitive if:
(A) for all $a \in A$, $(a, a) \in R$
(B) if $(a, b) \in R$, then $(b, a) \in R$
(C) if $(a, b) \in R$ and $(b, c) \in R$, then $(a, c) \in R$
(D) if $(a, b) \in R$ and $(b, a) \in R$, then $a = b$

197. The number of ways to choose r items from a set of n items with replacement, where order matters, is:
(A) P(n, r)
(B) n^r
(C) C(n, r)
(D) n*r

198. What distinguishes a ring from a group?
(A) A ring must be commutative, while a group need not be.
(B) A group must have an identity, while a ring need not.
(C) A ring involves integers, while a group involves real numbers.
(D) A ring has two binary operations, while a group has one.

199. Match the items in List I with those in List II:
List I (Graph Concept):
P. Eulerian Circuit
Q. Hamiltonian Circuit
R. Spanning Tree
S. Bipartite Graph.
List II (Property):
1. Visits every vertex exactly once
2. No odd-length cycles
3. Visits every edge exactly once
4. A connected acyclic subgraph with all

vertices.
(A) P-1, Q-3, R-4, S-2
(B) P-3, Q-1, R-2, S-4
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-4, R-1, S-3

200. Which of the following is a correct representation of a Boolean function?
(A) Flowchart
(B) PERT diagram
(C) Truth Table
(D) Pseudocode

201. The region of feasible solutions for a linear programming problem on a graph is:
(A) Always a convex polygon
(B) Always a concave polygon
(C) Sometimes convex, sometimes concave
(D) Always a circle

202. The term 'isomorphism' in the context of algebraic structures implies that two structures are:
(A) Fundamentally different but share limited properties.
(B) Subsets of one another in a larger set.
(C) Both identified as subgroups of a larger group.
(D) Structurally identical, just with different names for elements and operations.

203. The process of converting a propositional formula into an equivalent formula in Conjunctive Normal Form (CNF) is a fundamental step in:
(A) Proving tautologies using truth tables.
(B) Automated theorem proving using resolution.
(C) Simplifying Boolean circuits.
(D) Quantifier elimination.

204. Which of the following statements about relations is true?
(A) If a relation is reflexive, it cannot be anti-symmetric.
(B) If a relation is symmetric, it must also be transitive.
(C) The empty relation on a non-empty set is symmetric and transitive.
(D) A relation cannot be both symmetric and anti-symmetric.

205. In a conditional probability P(A|B), what does B represent?
(A) The event whose probability is to be found.
(B) The intersection of A and B.
(C) The sample space.
(D) The event that is known to have already occurred.

206. A path in a graph that does not repeat vertices is called a:
(A) Simple path
(B) Circuit
(C) Walk
(D) Closed walk

207. The canonical form of a Boolean expression is also known as:
(A) Simplified form
(B) Standard form
(C) Minimized form
(D) Complex form

208. In the context of project scheduling, what does cost consideration primarily refer to?
(A) The initial budget allocated for the project.
(B) The trade-off between project duration and project cost (crashing).
(C) The cost of resources used in the project.
(D) The penalty cost for late project completion.

209. Consider the following statements about algebraic structures:
(i) A group is a monoid in which every element has an inverse.
(ii) A ring is an abelian group under addition and a semi-group under multiplication, with distributivity holding.
(iii) An integral domain has no zero divisors.
(iv) A field is a commutative ring where every element has a multiplicative inverse.
(v) Every field is a ring, but not every ring is a field.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iv) only
(B) (i), (iii), (v) only
(C) (ii), (iii), (iv), (v) only
(D) All (i), (ii), (iii), (iv), and (v) are correct.

210. Which rule of inference allows one to conclude P from the premise $P \wedge Q$?
(A) Conjunction

Click Here to Buy book on Amazon

(B) Conjunction Introduction
(C) Simplification
(D) Modus Ponens

211. What does a rooted tree have that a general tree does not?
(A) Edges and vertices
(B) No cycles
(C) Leaves
(D) A distinguished vertex called the root

212. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The graphical solution of a linear programming problem is limited to problems with two or three decision variables.
Reason (R) : It is impossible to visualize the feasible region geometrically in more than three dimensions.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

213. Let R be a relation on the set of all lines in a plane defined by $(l1, l2) \in R$ if l1 is parallel to l2. This relation is:
(A) Only reflexive and symmetric
(B) Reflexive, symmetric, and transitive
(C) Only symmetric and transitive
(D) Only reflexive

214. The Inclusion-Exclusion Principle for three sets A, B, and C is given by $|A \cup B \cup C| = ?$
(A) $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C|$
(B) $|A| + |B| + |C| + |A \cap B| + |A \cap C| + |B \cap C| - |A \cap B \cap C|$
(C) $|A| + |B| + |C|$
(D) $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|$

215. In the context of graph theory, what is a multigraph?
(A) A graph where multiple edges (parallel edges) are allowed between the same pair of vertices.
(B) A graph with multiple disconnected components.
(C) A graph with weighted edges.
(D) A graph with directed edges.

216. The absorption law in Boolean algebra is represented by which of the following?
(A) $A + A' = 1$
(B) $A(A+B) = B$
(C) $A + AB = A$
(D) $A + BC = (A+B)(A+C)$

217. In a maximization LPP, if the simplex method terminates and one or more artificial variables are in the basis at a positive level, the problem has:
(A) An infeasible solution
(B) An optimal solution
(C) An unbounded solution
(D) Multiple optimal solutions

218. Consider the following statements:
Statement I : Every cyclic group is abelian.
Statement II : Every abelian group is cyclic.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

NODIA

219. The pre-order traversal of a binary tree follows which pattern?
(A) Left, Right, Root
(B) Root, Left, Right
(C) Left, Root, Right
(D) Root, Right, Left

220. Which of the following is not an application of predicate logic?
(A) Expressing mathematical statements
(B) Formalizing arguments in natural language
(C) Simplifying Boolean expressions
(D) Specifying properties in computer programs

221. An Equivalence relation partitions a set into disjoint subsets called:
(A) Partial orders
(B) Subgroups
(C) Kernels
(D) Equivalence classes

222. What is the minimum number of edges that must be removed from a connected graph with n vertices and m edges to form a spanning tree?

(A) m - n + 1
(B) m - n
(C) n - m + 1
(D) m + n

223. Which of the following is a characteristic of a PERT network diagram?
(A) It is activity-on-node.
(B) It is event-oriented.
(C) It allows for cycles.
(D) It uses deterministic time estimates.

224. A homomorphism 'f' from a group (G, *) to (H, •) is a function such that for all a, b in G:
(A) f(a * b) = f(a) * f(b)
(B) f(a * b) = f(a) • f(b)
(C) f(a * b) = f(b) • f(a)
(D) f(a • b) = f(a) * f(b)

225. Match the items in List I with those in List II regarding graph properties:
List I (Term):
P. Vertex Degree
Q. Path
R. Cycle
S. Simple Graph.
List II (Definition):
1. A path that starts and ends at the same vertex
2. A sequence of adjacent vertices
3. A graph with no loops or parallel edges
4. The number of edges incident to a vertex.
(A) P-4, Q-2, R-1, S-3
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-1, R-2, S-3
(D) P-2, Q-4, R-1, S-3

226. A Boolean expression consisting only of a sum of product terms (SOP) is in what form?
(A) Conjunctive Normal Form
(B) Product of Maxterms
(C) Disjunctive Normal Form
(D) Propositional Form

227. What is the primary advantage of the dual simplex method?
(A) It is faster than the regular simplex method.
(B) It can handle integer variables.
(C) It works for non-linear problems.
(D) It is efficient for sensitivity analysis when a new constraint is added.

NODIA

228. Consider the following statements about Linear Programming:
(i) The objective function must be linear.
(ii) All constraints must be linear equalities or inequalities.
(iii) Decision variables must be non-negative.
(iv) The optimal solution must occur at a vertex of the feasible region.
(v) It can be used for project scheduling without modification.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iii) only
(B) (i), (ii), (iii), (iv) only
(C) All (i), (ii), (iii), (iv), (v) are true
(D) (i), (ii), (iv) only

229. A proposition is a statement that is:
(A) Always true.
(B) Either true or false, but not both.
(C) Always false.
(D) A question or command.

230. If a relation is reflexive, symmetric, and transitive, it is called:
(A) A partial ordering
(B) A total ordering
(C) A function
(D) An equivalence relation

231. What is the purpose of the inductive step in mathematical induction?
(A) To show that if the statement holds for an arbitrary case k, it must also hold for the next case k+1.
(B) To prove the base case P(1) is true.
(C) To assume the conclusion is true.
(D) To check for counterexamples.

232. The set of all invertible n x n matrices with matrix multiplication forms a:
(A) Ring
(B) Field
(C) Group
(D) Semi-group only

233. In a post-order traversal of a binary tree, the root is visited:
(A) Last
(B) First
(C) Second
(D) Before the right subtree

234. The purpose of simplifying a Boolean function is to:
(A) Increase the propagation delay.
(B) Make the function more complex.
(C) Convert it to a POS form.
(D) Obtain a logic circuit with the minimum number of gates.

235. The assignment model is a special case of the transportation model where:
(A) The number of sources equals the number of destinations.
(B) All supply and demand values are equal to 1.
(C) The cost matrix is symmetric.
(D) The objective is to maximize profit.

236. The contrapositive of $p \rightarrow q$ is:
(A) $q \rightarrow p$
(B) $\neg p \rightarrow \neg q$
(C) $\neg q \rightarrow \neg p$
(D) $p \wedge \neg q$

237. A partially ordered set in which every pair of elements has a least upper bound and a greatest lower bound is called a:
(A) Semilattice
(B) Poset
(C) Chain
(D) Lattice

238. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A connected graph with n vertices and n-1 edges is a tree.
Reason (R) : A graph is a tree if and only if it is connected and acyclic.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

239. Consider the following statements:
Statement I : The identity element in a group is unique.
Statement II : The inverse of each element in a group is unique.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

240. What is the main characteristic of a critical path in a project network?
(A) It is the shortest path.
(B) It has zero slack.
(C) It involves the least expensive activities.
(D) It can be delayed without affecting the project finish time.

241. The number of spanning trees in a complete graph K4 is:
(A) 16
(B) 4
(C) 8
(D) 24

242. The expression $\neg(p \vee q)$ is logically equivalent to:
(A) $\neg p \vee \neg q$
(B) $p \wedge q$
(C) $\neg p \wedge \neg q$
(D) $\neg p \vee q$

243. In Bayes' Theorem, $P(A|B) = [P(B|A) * P(A)] / P(B)$, what is P(A) called?
(A) Posterior probability
(B) Likelihood
(C) Evidence
(D) Prior probability

244. A minimal spanning tree of a weighted graph is a spanning tree:
(A) with the maximum possible total edge weight.
(B) with the fewest number of edges.
(C) with the minimum possible total edge weight.
(D) that is also a complete graph.

245. Arrange the following events in the evolution of project management techniques in chronological order:
1. Development of CPM (Critical Path Method)
2. Development of PERT (Program Evaluation and Review Technique)
3. Use of Gantt Charts.
(A) 1, 2, 3
(B) 3, 2, 1

Downloaded from www.pdf.tube

(C) 2, 1, 3
(D) 3, 1, 2

246. A ring (R, +, *) is an integral domain if it is commutative, has a multiplicative identity, and:
(A) is finite
(B) every element has an inverse
(C) is a field
(D) has no zero divisors

247. The solution to an integer programming problem is always:
(A) An integer point on or within the feasible region of the corresponding LP problem.
(B) A point within the feasible region of the corresponding LP problem.
(C) The same as the rounded solution of the corresponding LP problem.
(D) A corner point of the feasible region of the corresponding LP problem.

248. Which of the following statements about graph theory is INCORRECT?
(A) A simple graph contains no loops.
(B) The sum of the degrees of all vertices is an even number.
(C) Every planar graph can be colored with 6 colors.
(D) A graph with n vertices must have at least n-1 edges to be connected.

249. Match the items in List I with those in List II:
List I:
P. Propositional Logic
Q. Predicate Logic
R. Set Theory
S. Group Theory.
List II:
1. Deals with objects and their properties using quantifiers (∀, ∃)
2. Studies algebraic structures with a single binary operation satisfying certain axioms
3. Foundation based on collections of objects
4. Deals with statements and their truth values using connectives (∧, ∨, →).
(A) P-4, Q-1, R-3, S-2
(B) P-1, Q-4, R-2, S-3
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-1, R-3, S-4

Direction : Please read the following source carefully.

Mathematical logic serves as the foundation for systematic reasoning in computer science and mathematics. It is broadly divided into propositional logic and predicate logic. Propositional logic deals with propositions, which are declarative statements that are either true or false, and combines them using logical connectives such as AND, OR, NOT, and IMPLIES. An important aspect of propositional logic is the determination of propositional equivalences, which identifies when two different compound propositions have the same truth values for all possible inputs. For standardization and simplification, logical expressions are often converted into normal forms, such as Conjunctive Normal Form (CNF) or Disjunctive Normal Form (DNF). Predicate logic extends propositional logic by introducing predicates, variables, and quantifiers. Predicates are properties or relationships that subjects of a statement can have. Quantifiers, namely the universal quantifier (∀, for all) and the existential quantifier (∃, there exists), allow for statements about entire classes of objects. The complexity increases with nested quantifiers, where the order of quantification significantly impacts the meaning. To derive conclusions from a set of premises, rules of inference, such as Modus Ponens and Modus Tollens, provide a structured method for constructing valid logical arguments.

Based on the above source, attempt the following five questions.

250. According to the passage, what is the primary purpose of converting logical expressions into normal forms?
(A) To introduce additional complexity for advanced analysis.
(B) For standardization and simplification of the expressions.
(C) To allow the use of nested quantifiers in the expression.
(D) To translate predicate logic into propositional logic.

251. Which of the following best describes the distinction between propositional logic and predicate logic as outlined in the text?
(A) Propositional logic deals with simple true/false statements, while predicate logic

NODIA

incorporates variables and quantifiers.
(B) Propositional logic uses rules of inference, whereas predicate logic uses normal forms.
(C) Propositional logic is used for mathematical proofs, while predicate logic is used for computer programming.
(D) Propositional logic uses connectives like AND and OR, while predicate logic exclusively uses IMPLIES.

252. Based on the source, what can be inferred about the role of 'rules of inference'?
(A) They are used to determine the truth value of a single proposition.
(B) They are exclusively used for simplifying expressions into CNF or DNF.
(C) They provide a mechanism for generating valid conclusions from given statements.
(D) They are tools for checking propositional equivalences between two expressions.

253. The passage mentions that the order of nested quantifiers significantly impacts the meaning. Which of the following hypothetical statements best illustrates this concept?
(A) A statement using only universal quantifiers ($\forall x\ \forall y$) has a fixed meaning regardless of variable order.
(B) A statement without any quantifiers has a more ambiguous meaning than one with nested quantifiers.
(C) A statement with only existential quantifiers ($\exists x\ \exists y$) is logically equivalent to one with universal quantifiers.
(D) A statement with mixed quantifiers, like $\forall x\ \exists y$, may have a different meaning than $\exists y\ \forall x$.

254. What are the fundamental components that predicate logic adds to the framework of propositional logic?
(A) Logical connectives, truth tables, and normal forms.
(B) Modus Ponens, Modus Tollens, and other rules of inference.
(C) Predicates, variables, and quantifiers.
(D) The concepts of true, false, and propositional equivalence.

Direction : Please read the following source carefully.

In discrete mathematics, sets and relations are fundamental structures. A set is a well-defined collection of distinct objects, and foundational operations on sets include union, intersection, difference, and complement. Relations extend this concept by defining a relationship between elements, often between elements of two sets or within a single set. A relation can be formally represented as a set of ordered pairs or visually through a matrix or a directed graph. The utility of relations is largely determined by their properties, such as being reflexive (every element is related to itself), symmetric (if a is related to b, then b is related to a), transitive (if a is related to b and b is related to c, then a is related to c), and antisymmetric (if a is related to b and b is related to a, then a must equal b). These properties are used to classify relations into significant types. An equivalence relation is a relation that is reflexive, symmetric, and transitive. Its key feature is that it partitions the underlying set into disjoint subsets known as equivalence classes. In contrast, a partially ordering relation, or a poset, is a relation that is reflexive, antisymmetric, and transitive. It imposes a hierarchical order on the elements of a set, but unlike a total order, it is not required that every pair of elements be comparable.

Based on the above source, attempt the following five questions.

255. A relation R on a set A is defined as an equivalence relation. Based on the passage, which combination of properties must R satisfy?
(A) Reflexive, Antisymmetric, Transitive.
(B) Reflexive, Symmetric, Transitive.
(C) Symmetric, Transitive, Antisymmetric.
(D) Reflexive, Symmetric, Comparable.

256. What is the primary distinguishing feature between a partially ordering relation and an equivalence relation as described in the text?
(A) An equivalence relation must be reflexive, while a partial order need not be.
(B) An equivalence relation partitions a set, while a partial order combines elements into a single class.
(C) A partial order requires all elements to be related, while an equivalence relation does not.
(D) A partial order uses the antisymmetric property, whereas an equivalence relation

Click Here to Buy book on Amazon

uses the symmetric property.

257. According to the source, what is a key consequence of applying an equivalence relation to a set?
(A) It divides the set into non-overlapping equivalence classes.
(B) It sorts the set into a linear, hierarchical order.
(C) It reduces the number of elements in the set.
(D) It proves that all elements in the set are identical.

258. If a relation on a set of integers is defined as a is related to b if a ? b, which property mentioned in the passage would prevent it from being an equivalence relation?
(A) The reflexive property, because an integer is not always less than or equal to itself.
(B) The transitive property, because if $a \leq b$ and $b \leq c$, it does not imply $a \leq c$.
(C) The symmetric property, because if $a \leq b$, it does not necessarily mean $b \leq a$.
(D) The antisymmetric property, because it is required for equivalence relations.

259. The passage states that in a partially ordering, it is not required that every pair of elements be comparable. What does this imply?
(A) There can exist elements 'a' and 'b' in the set for which neither 'a is related to b' nor 'b is related to a' is true.
(B) The relation is not transitive for all pairs of elements.
(C) The relation might not be reflexive for every element.
(D) The set cannot be partitioned into equivalence classes.

Direction : Please read the following source carefully.

Combinatorics and discrete probability provide essential tools for counting and analyzing uncertain events. The basics of counting are founded on principles like the sum and product rules. A more nuanced counting argument is the Pigeonhole Principle, which asserts that if 'n' items are put into 'm' containers, with n > m, then at least one container must contain more than one item. When arranging objects, a distinction is made between permutations, which are ordered arrangements, and combinations, which are unordered selections. For more complex counting scenarios involving overlapping sets, the Inclusion-Exclusion Principle provides a systematic way to count the number of elements in the union of multiple sets by adding the sizes of the individual sets and subtracting the sizes of their intersections. Mathematical induction is a powerful proof technique used to establish that a given statement is true for all natural numbers. It involves proving a base case and then showing that if the statement holds for an arbitrary case 'k', it must also hold for the case 'k+1'. In the realm of uncertainty, discrete probability assigns a probability to each outcome in a finite or countably infinite sample space. A cornerstone of conditional probability is Bayes' Theorem, which describes the probability of an event based on prior knowledge of conditions that might be related to the event. It provides a way to update probability estimates when new evidence becomes available.

Based on the above source, attempt the following five questions.

NODIA

260. A university department has 10 professors and needs to form a 3-member committee. The order in which the professors are chosen does not matter. Which counting technique described in the passage is appropriate for this task?
(A) Permutations.
(B) The Pigeonhole Principle.
(C) The Inclusion-Exclusion Principle.
(D) Combinations.

261. What is the fundamental purpose of Bayes' Theorem as explained in the source?
(A) To calculate the total number of outcomes in a sample space.
(B) To revise the probability of an event given new, related information.
(C) To prove mathematical statements for all natural numbers.
(D) To count the number of items in overlapping sets.

262. According to the passage, the principle of mathematical induction requires the completion of which two steps?
(A) Defining a permutation and a combination.
(B) Applying the sum rule and the product rule.

(C) Proving a base case and an inductive step.
(D) Using the Inclusion-Exclusion Principle and Bayes' Theorem.

263. A survey finds that in a group of students, 50 take History, 40 take Geography, and 20 take both. To find the total number of students who take at least one of these subjects, which principle would be most directly applied?
(A) The Pigeonhole Principle.
(B) Mathematical Induction.
(C) Bayes' Theorem.
(D) The Inclusion-Exclusion Principle.

264. If you have 5 email accounts but you receive 6 distinct emails, what can you conclude based on the Pigeonhole Principle?
(A) At least one email account must have received more than one email.
(B) Each email account has received at least one email.
(C) One specific email account must have received exactly two emails.
(D) It is impossible to determine the distribution of emails.

Direction : Please read the following source carefully.

Group theory is a branch of abstract algebra that studies algebraic structures known as groups. An algebraic structure consists of a set and one or more binary operations. A semi-group is a set with an associative binary operation. A structure is elevated to a group if it satisfies four axioms: closure under the operation, associativity of the operation, the existence of an identity element, and the existence of an inverse element for each element. Subgroups are subsets of a group that themselves form a group under the same operation. Further complexity is explored through product and quotients of algebraic structures. The relationships between different algebraic structures are formalized through mappings. A homomorphism is a map between two algebraic structures that preserves the operation. If a homomorphism is also a bijection (both one-to-one and onto), it is called an isomorphism, indicating the two structures are algebraically identical. An automorphism is a special case of an isomorphism where the mapping is from a group to itself. Beyond groups, other structures are defined. A ring is a set with two binary operations (typically addition and multiplication) that satisfies certain properties. An integral domain is a commutative ring that has no zero-divisors. A field is a commutative ring in which every non-zero element has a multiplicative inverse, making it a more versatile structure for arithmetic.

Based on the above source, attempt the following five questions.

265. What additional property must a semi-group possess to be classified as a group?
(A) It must be commutative.
(B) It must have an identity element and an inverse for each element.
(C) It must have two binary operations instead of one.
(D) It must be a finite set.

266. Based on the passage, what is the key difference between a homomorphism and an isomorphism?
(A) A homomorphism preserves the operation, while an isomorphism does not.
(B) An isomorphism is a bijective mapping, while a homomorphism is not necessarily bijective.
(C) A homomorphism maps a group to itself, while an isomorphism maps between different groups.
(D) An isomorphism applies to rings and fields, while a homomorphism only applies to groups.

267. An automorphism is described as a specific type of isomorphism. What is its defining characteristic?
(A) It is a mapping from a group onto itself.
(B) It maps a group to a completely different algebraic structure.
(C) It does not require the preservation of the group operation.
(D) It only applies to non-commutative groups.

268. How does the source distinguish a field from an integral domain?
(A) A field must be commutative, while an integral domain need not be.
(B) A field has no zero-divisors, which is a property not found in integral domains.
(C) Every non-zero element in a field has a multiplicative inverse, a condition not

Downloaded from www.pdf.tube

required for an integral domain.

(D) An integral domain has two binary operations, whereas a field has only one.

269. A subset H of a group G is considered a subgroup. What condition must H satisfy according to the text?

(A) H must be larger in size than G.

(B) H must be an empty set.

(C) H must have a different binary operation than G.

(D) H must form a group under the same operation as G.

Direction : Please read the following source carefully.

Graph theory involves the study of graphs as mathematical representations of networks. A fundamental distinction exists between a simple graph, which has no loops or multiple edges between the same two vertices, and a multigraph, which allows for multiple edges. When edges are assigned a numerical value, such as cost or distance, the graph is known as a weighted graph. Navigating through a graph is described by paths and circuits. A path is a sequence of edges connecting a sequence of vertices, while a circuit is a path that starts and ends at the same vertex. A central problem in weighted graphs is finding the shortest paths between vertices. Special types of circuits are of particular interest. An Eulerian circuit is a circuit that traverses every edge of the graph exactly once. In contrast, a Hamiltonian circuit is a circuit that visits every vertex in the graph exactly once (except for the start/end vertex). The existence of such circuits depends on the properties of the graph, such as the degrees of its vertices. An Eulerian path is a path that traverses every edge exactly once but starts and ends at different vertices. Similarly, a Hamiltonian path visits every vertex exactly once but does not form a closed loop. These concepts are foundational to solving problems in logistics, network routing, and circuit design.

Based on the above source, attempt the following five questions.

270. What is the primary characteristic that distinguishes a weighted graph from a simple graph?

(A) A weighted graph must contain a Hamiltonian circuit.

(B) A weighted graph allows for multiple edges between two vertices.

(C) A weighted graph has numerical values associated with its edges.

(D) A weighted graph cannot have any circuits.

271. According to the passage, what is the core difference between an Eulerian circuit and a Hamiltonian circuit?

(A) An Eulerian circuit is for weighted graphs, while a Hamiltonian circuit is for unweighted graphs.

(B) An Eulerian circuit visits every edge once, while a Hamiltonian circuit visits every vertex once.

(C) An Eulerian circuit starts and ends at different vertices, while a Hamiltonian circuit does not.

(D) An Eulerian circuit is a path, while a Hamiltonian circuit is a more complex structure.

272. Which of the following scenarios best represents the problem of finding a Hamiltonian path?

(A) A mail carrier needing to travel along every street in a neighborhood exactly once to deliver mail.

(B) Finding the cheapest route to fly from New York to Los Angeles.

(C) A network administrator checking the connectivity of all data links in a network.

(D) A salesperson wanting to visit every city on their route exactly once before returning home.

NODIA

273. Based on the definitions provided, what differentiates a path from a circuit?

(A) A circuit is a path that begins and ends at the same vertex.

(B) A path connects vertices, while a circuit connects edges.

(C) A path must not repeat vertices, while a circuit must.

(D) A circuit can only exist in a multigraph, not a simple graph.

274. The problem of finding the shortest paths in weighted graphs implies an optimization goal. What is being optimized?

(A) The number of vertices visited.

(B) The number of edges in the path.
(C) The sum of the weights of the edges along the path.
(D) The number of circuits in the graph.

Direction : Please read the following source carefully.

Advanced graph theory explores structural properties and applications, including planarity, coloring, and tree structures. A planar graph is one that can be embedded in a plane such that no two edges cross each other. This property is crucial in areas like circuit board design. Closely related is the concept of graph coloring, which is an assignment of labels, traditionally called colors, to the vertices of a graph such that no two adjacent vertices share the same color. A significant subclass of colorable graphs is the set of bipartite graphs, which are graphs whose vertices can be divided into two disjoint and independent sets, U and V, such that every edge connects a vertex in U to one in V. Essentially, bipartite graphs are 2-colorable. Trees represent a particularly important class of graphs; they are connected graphs that contain no cycles. A rooted tree is a tree in which one vertex has been designated as the root, which imposes a hierarchical structure. Tree traversals are algorithms for visiting each node in a tree exactly once, with common methods being pre-order, in-order, and post-order. For any connected graph, a spanning tree is a subgraph that is a tree and includes all the vertices of the original graph. Related to this is a cut-set, which is a set of edges whose removal increases the number of connected components of a graph. Trees are also applied in data compression, where they can be used to generate prefix codes, a type of code where no codeword is a prefix of any other codeword.

Based on the above source, attempt the following five questions.

275. According to the passage, what is the defining property of a planar graph?
(A) It can be drawn on a flat surface without any edges intersecting.
(B) It must be a complete graph where all vertices are connected.
(C) It must be a tree with no cycles.
(D) All of its vertices must have the same color.

276. The source defines a bipartite graph as being 2-colorable. What does this imply about the connections within a bipartite graph?
(A) Every vertex is connected to every other vertex.
(B) The graph is divided into two sets, and edges only connect vertices within the same set.
(C) The graph must contain at least one triangle (a cycle of length 3).
(D) The graph is divided into two sets, and edges only connect vertices from one set to the other.

277. Which of the following best describes the relationship between a connected graph G and its spanning tree T?
(A) T contains all the edges of G but only some of the vertices.
(B) T is a subgraph of G that is a tree and connects all of G's vertices.
(C) T is always a planar graph, even if G is not.
(D) T is formed by adding edges to G to remove all cycles.

278. What is the main objective of graph coloring as described in the text?
(A) To ensure that every vertex has a unique color.
(B) To find the shortest path between two vertices of the same color.
(C) To assign colors to vertices such that no adjacent vertices have the same color.
(D) To identify the root of the graph.

279. The passage mentions that trees can be used to generate prefix codes. What is the stated characteristic of a prefix code?
(A) All codewords must have the same length.
(B) Every codeword must start with the same symbol or prefix.
(C) Codewords are generated using a post-order traversal exclusively.
(D) No codeword can be the beginning part of another codeword.

Direction : Please read the following source carefully.

Boolean algebra is the mathematical foundation for digital logic design and computer architecture. It is a system of algebra that operates on binary variables, typically represented as 0 (false)

Click Here to Buy book on Amazon

and 1 (true), using logical operations like AND (conjunction), OR (disjunction), and NOT (negation). Any logical expression that can be evaluated to a binary value is known as a Boolean function. These functions can be formally described or represented in several ways. The most fundamental representation is a truth table, which exhaustively lists the output of the function for all possible combinations of input variables. For algebraic manipulation, Boolean functions are expressed as expressions. These expressions can be standardized into specific formats, such as the Sum of Products (SOP) form, where the function is a sum of product terms, or the Product of Sums (POS) form, where it is a product of sum terms. Canonical forms, which are unique for any given function, include the canonical SOP (using minterms) and canonical POS (using maxterms). A primary goal in digital circuit design is the simplification of Boolean functions. This process aims to find an equivalent function that requires the minimum number of literals (variables or their complements) or terms, which directly translates to a less complex and more economical logic circuit implementation, without altering the input-output relationship defined by the original function.

Based on the above source, attempt the following five questions.

280. What is the primary motivation for the simplification of Boolean functions as stated in the passage?
(A) To reduce the cost and complexity of the corresponding logic circuit.
(B) To change the function's output for specific inputs.
(C) To make the function compatible with non-binary systems.
(D) To convert the function into a truth table representation.

281. The source describes a truth table as a form of representation for a Boolean function. What information does a truth table provide?
(A) The simplest algebraic expression for the function.
(B) The function's output for every possible combination of its inputs.
(C) The number of logic gates required for implementation.
(D) The canonical POS form of the function using maxterms.

282. Which of the following best describes the Sum of Products (SOP) form?
(A) An expression formed by multiplying several sum terms together.
(B) An expression formed by adding several product terms together.
(C) An expression that uses only the NOT and AND operations.
(D) An expression represented exclusively by a truth table.

283. What are the fundamental elements that Boolean algebra operates on, according to the text?
(A) Binary variables and logical operations.
(B) Real numbers and arithmetic operations.
(C) Sets and relational operators.
(D) Graphs and their vertices.

284. The passage refers to minterms and maxterms in the context of canonical forms. This implies that canonical forms are:
(A) A simplified but non-unique representation of a function.
(B) A graphical representation of the Boolean function.
(C) A standardized representation that is unique for any given function.
(D) An optional step in circuit design that increases complexity.

Direction : Please read the following source carefully.

Optimization through Linear Programming (LP) is a powerful technique for achieving the best outcome, such as maximum profit or lowest cost, in a mathematical model whose requirements are represented by linear relationships. The formulation of an LP problem requires creating a mathematical model consisting of three key components: decision variables, which represent the quantities to be determined; an objective function, which is a linear expression of the decision variables to be maximized or minimized; and constraints, which are a set of linear equalities or inequalities that limit the values of the decision variables. For problems involving only two decision variables, a graphical solution

Downloaded from www.pdf.tube

method can be employed. This involves plotting the constraints to identify a feasible region, and the optimal solution is guaranteed to lie at one of the corner points of this region. For problems with more variables, the Simplex method is an efficient, iterative algorithm. It starts at a vertex of the feasible region and systematically moves to adjacent vertices, each time improving the value of the objective function until an optimal solution is reached. The Dual Simplex method is a variant useful for certain types of problems and for post-optimality analysis. An important part of the optimization process is sensitivity analysis, which investigates how the optimal solution is affected by changes in the model's parameters, such as the coefficients of the objective function or the right-hand-side values of the constraints.

Based on the above source, attempt the following five questions.

285. According to the passage, what are the three essential components of a Linear Programming mathematical model?
(A) Graphical solution, Simplex method, and sensitivity analysis.
(B) Vertices, feasible region, and corner points.
(C) Maximum profit, lowest cost, and optimal outcome.
(D) Decision variables, objective function, and constraints.

286. The graphical solution method is mentioned as being suitable for a specific type of LP problem. What is this type?
(A) Problems with a non-linear objective function.
(B) Problems with an undefined feasible region.
(C) Problems with exactly two decision variables.
(D) Problems involving integer variables only.

287. What is the fundamental procedure of the Simplex method as described in the source?
(A) It randomly tests points within the feasible region.
(B) It moves iteratively between vertices of the feasible region to improve the objective function.
(C) It graphically plots the constraints to visually identify the best solution.
(D) It analyzes the sensitivity of the solution to parameter changes.

288. What is the purpose of sensitivity analysis in the context of Linear Programming?
(A) To formulate the initial mathematical model.
(B) To find the initial feasible solution for the Simplex method.
(C) To convert a maximization problem into a minimization problem.
(D) To determine how the optimal solution changes when model parameters are altered.

289. In the graphical method for LP, where is the optimal solution guaranteed to be located?
(A) At a corner point of the feasible region.
(B) At the geometric center of the feasible region.
(C) Outside the feasible region, closest to the origin.
(D) Along the edge connecting the first two constraints.

Direction : Please read the following source carefully.

Within the broader field of mathematical optimization, several specialized models address specific problem structures. Integer Programming (IP) is a significant extension of linear programming. It is applied when decision variables must take on integer values, as fractional results are often meaningless in real-world scenarios like assigning people to tasks or manufacturing discrete units. The constraint that variables must be integers makes IP problems computationally more complex than standard LP problems. The Transportation Model is a classic type of optimization problem that seeks to minimize the cost of shipping goods from a set of sources to a set of destinations. This model is characterized by supply constraints at each source and demand constraints at each destination. The goal is to determine the quantity of goods to be shipped from each source to each destination to satisfy all demands while not exceeding supplies, all at the minimum total cost. A special, more constrained version of the transportation problem is the Assignment Model. Its primary objective is to assign a number of resources (e.g., employees) to an equal number of tasks on a one-to-one basis. The model aims to find the assignment

NODIA

that minimizes the total cost or time. The defining feature of the assignment model is that the supply at each source and the demand at each destination are all exactly equal to one.

Based on the above source, attempt the following five questions.

290. What is the primary characteristic that distinguishes Integer Programming (IP) from standard Linear Programming (LP)?
(A) IP problems can only be minimized, not maximized.
(B) IP problems do not use an objective function.
(C) IP problems involve decision variables that are restricted to be integers.
(D) IP problems are computationally less complex than LP problems.

291. According to the passage, what is the main objective of the Transportation Model?
(A) To minimize the cost of distributing goods from sources to destinations.
(B) To maximize the number of routes between sources and destinations.
(C) To assign each source to exactly one destination.
(D) To ensure that supply at each source is exactly one.

292. The passage describes the Assignment Model as a special case of the Transportation Model. What is its key distinguishing feature?
(A) The number of sources must be greater than the number of destinations.
(B) To maximize the total cost of assignment.
(C) The decision variables are allowed to be fractional.
(D) All supply and demand values in the model are equal to one.

293. In which of the following scenarios would an Integer Programming model be more appropriate than a standard Linear Programming model?
(A) Determining the optimal blend of different oils to produce a final product.
(B) Deciding the number of full-time and part-time staff to hire.
(C) Calculating the amount of a chemical to produce, measured in liters.
(D) Optimizing the allocation of investment capital across different stocks.

294. Both the Transportation and Assignment models are described as aiming to minimize cost. What are the primary constraints in a Transportation Model?
(A) Time limits and resource availability.
(B) One-to-one assignment rules.
(C) Supply limits at sources and demand requirements at destinations.
(D) Integer-only decision variables.

Direction : Please read the following source carefully.

For managing large-scale projects, PERT (Program Evaluation and Review Technique) and CPM (Critical Path Method) are indispensable optimization tools. They are used for the planning, scheduling, and control of complex activities with interdependencies. Both methodologies rely on a diagram representation, usually a network diagram, where activities are shown as nodes or arcs, illustrating the project's workflow and dependencies. A central element of these methods is the Critical Path Calculation. The critical path is defined as the sequence of activities that forms the longest path through the network diagram. The length of this path determines the shortest possible duration for the entire project. Any delay in an activity on the critical path will directly delay the project's completion. Activities on this path are characterized by having zero slack or float. Beyond just scheduling, these methods incorporate broader project management considerations. Resource levelling is a technique used to balance the demand for specific resources throughout the project's lifecycle. It aims to avoid periods of underutilization or excessive demand by adjusting the schedule of non-critical activities. Furthermore, cost consideration in project scheduling is a vital aspect, often involving an analysis known as crashing. This involves evaluating the trade-offs between allocating more resources to shorten an activity's duration (and thus the project's duration) and the increased costs associated with this acceleration.

Based on the above source, attempt the following five questions.

295. What defines the critical path in a project network diagram according to the passage?
(A) The path with the fewest number of activities.

Click Here to Buy book on Amazon

(B) The shortest possible path from the start to the end of the project.
(C) The sequence of activities that has the most available resources.
(D) The longest path through the network, determining the minimum project duration.

296. Based on the description, what is the significance of an activity being on the critical path?
(A) Any delay in this activity will lead to a delay in the overall project.
(B) It can be delayed without affecting the project's completion time.
(C) It requires the fewest resources to complete.
(D) It is an optional activity that can be skipped if necessary.

297. What is the primary goal of the resource levelling technique mentioned in the source?
(A) To calculate the critical path of the project.
(B) To balance the utilization of resources over the project's duration.
(C) To reduce the total cost of the project to its absolute minimum.
(D) To identify the activities with the highest risk of failure.

NODIA

298. The passage discusses cost consideration in project scheduling, including crashing. What trade-off is being evaluated in this process?
(A) The trade-off between project quality and project speed.
(B) The trade-off between increasing costs and decreasing project duration.
(C) The trade-off between using PERT and using CPM.
(D) The trade-off between resource levelling and resource allocation.

299. PERT and CPM are described as tools for planning, scheduling, and control. What is the foundational component that both methods use to model a project?
(A) A network diagram representing activities and dependencies.
(B) A detailed financial budget.
(C) A list of project personnel.
(D) A set of quality control standards.

ANSWERS

1. (B) A proposition
In logic, a proposition is defined as a declarative statement that carries a specific truth value, being either true or false. Unlike commands, questions, or exclamations, which do not assert a fact, a proposition establishes a clear claim about the state of affairs. This property allows propositions to be used as the fundamental building blocks of logical reasoning.

2. (C) Statement I is true, but Statement II is false.
The union of two sets includes all unique elements present in either set or both. Conversely, the intersection represents only those elements shared by both sets simultaneously. Statement I correctly describes the union operation. However, Statement II is inaccurate because the intersection is defined by elements common to both sets, not elements present in one and absent from the other.

3. (C) If n pigeons are put into m pigeonholes with $n > m$, then at least one pigeonhole must contain more than one pigeon.
The basic version of this principle asserts that if the number of items being distributed exceeds the number of available containers, at least one container must hold multiple items. This mathematical concept is widely used in combinatorics to prove the existence of specific configurations or overlaps within finite sets where the allocation of elements exceeds the capacity of the target set.

4. (D) Closure and Associativity
For an algebraic structure to be classified as a semi-group, it must fulfill two primary requirements under its binary operation. First, it must satisfy the closure property, ensuring that the operation on any two elements results in another element within the set. Second, the operation must be associative, meaning the grouping of elements does not change the final result obtained.

5. (C) A circuit
Within the field of graph theory, a circuit is defined as a closed walk where the initial and terminal vertices are identical. While it allows for the repetition of edges and vertices, the essential characteristic is that the path forms a complete loop. This concept is distinct from a simple path or a trail, which focuses on different edge-vertex restrictions.

6. (B) $A + B$
Simplification of this Boolean expression relies on the distributive law, where A plus the product of its complement and B is equivalent to A plus B. In logic design, this simplification reduces the number of required gates and inputs. By applying algebraic identities, the term involving the complement is absorbed, resulting in a more efficient expression used in digital circuit optimization.

7. (D) The objective function must be a quadratic equation.
A standard linear programming model requires a linear objective function and a set of linear constraints to determine an optimal outcome. Decision variables are traditionally restricted to non-negative values to reflect real-world limitations. Requiring the objective function to be quadratic would categorize the problem as a quadratic programming model instead, which involves different mathematical properties and solving techniques than linear models.

8. (A) The expression $p \rightarrow q$ is a tautology.
A tautology is a statement that is true under every possible truth value assignment of its components. The implication p implies q is not a tautology because it becomes false when the antecedent is true and the consequent is false. In contrast, the law of excluded middle and conjunction with a negation represent standard examples of tautological and contradictory logical forms respectively.

9. (C) Reflexive, Symmetric, and Transitive
For a relation on a set to be categorized as an equivalence relation, it must concurrently satisfy three fundamental properties. It must be reflexive, meaning every element relates to itself.

Downloaded from www.pdf.tube

It must be symmetric, where if one element relates to another, the reverse is also true. Finally, it must be transitive, ensuring that relationships propagate through intermediate elements within the set.

10. (A) 120
The word LOGIC consists of five distinct letters with no repetitions. To find the total number of unique arrangements, one calculates the factorial of the number of letters. Since there are five positions to fill, the first has five choices, the second has four, and so on. Multiplying these values results in one hundred and twenty distinct possible linear permutations.

11. (D) Commutativity
An Abelian group is a specific type of group that satisfies the standard group axioms while also adhering to the commutative property. This means that for any two elements within the group, the result of the binary operation remains unchanged regardless of the order in which the elements are combined. This additional symmetry distinguishes Abelian groups from more general, non-commutative group structures.

12. (B) Simple graph
A simple graph is defined by the absence of both self-loops and parallel edges between its vertices. This restriction ensures that any pair of vertices is connected by at most one edge and that no edge connects a vertex to itself. Such graphs are fundamental in modeling relationships where only distinct connections between different entities are relevant for the mathematical analysis.

13. (A) Both Statement I and Statement II are true.
Boolean functions are mathematical expressions that map binary inputs to a binary output, and their behavior can be fully detailed using a truth table. Furthermore, Karnaugh maps provide a visual and systematic method for simplifying these functions by identifying adjacent terms that can be combined. Both of these statements correctly identify foundational tools used in the analysis and design of digital systems.

14. (D) Two variables
The graphical method is an optimization technique used to solve linear programming problems by plotting constraints on a coordinate plane. This approach is practically limited to problems containing exactly two decision variables because they can be easily represented on a two-dimensional axis. Higher-dimensional problems require more complex algebraic methods, such as the Simplex algorithm, as they cannot be visualized through simple geometry.

15. (A) If the ground is not wet, then it is not raining.
The contrapositive of a logical implication is formed by negating both the antecedent and the consequent and then interchanging them. If the original statement is represented as if p then q, its contrapositive is if not q then not p. In logic, an implication and its contrapositive are always logically equivalent, meaning they share the exact same truth table values.

16. (B) An Equivalence Relation
This relation on integers is reflexive since zero is always even. It is symmetric because if the difference between two numbers is even, the difference in reverse order is also even. Finally, it is transitive because the sum of two even differences is also even. Since all three properties are satisfied, the relation qualifies as an equivalence relation partitioning the integers.

17. (B) 30
To find the number of students who play at least one sport, the principle of inclusion-exclusion is applied. By adding the number of cricket players and football players, and then subtracting those who play both, the total unique count is determined. In this specific scenario, the sum of fifteen and twenty minus five results in exactly thirty students participating in sports.

18. (A) An isomorphism
An isomorphism in group theory is a function between two groups that preserves the group operation while being both injective and surjective. This means every element in the first group maps to exactly one unique element in the second group, and vice versa. Such a mapping demonstrates that the two groups are structurally identical, regardless of the specific nature of their individual elements.

19. (C) An Eulerian path

NODIA

An Eulerian path is a specific type of traversal in graph theory that requires visiting every edge exactly once. This concept is distinct from Hamiltonian paths, which focus on visiting every vertex. For such a path to exist in a connected graph, the graph must have either zero or exactly two vertices with an odd degree, according to Euler's fundamental theorem.

20. (D) Tautology
In propositional logic and Boolean algebra, a tautology refers to a statement or function that yields a true or high output for every possible combination of input values. It represents a logical certainty that does not depend on the specific truth values of its constituent variables. This property makes tautologies useful for establishing fundamental logical equivalences and rules of inference during proofs.

21. (C) To convert 'less than or equal to' inequalities into equalities.
Slack variables are introduced in the Simplex method to transform linear inequalities of the less than or equal to type into equivalent linear equations. By adding a non-negative variable to the left side of the inequality, the difference between the actual usage and the available limit is represented. This conversion is necessary to perform the matrix operations required by the algorithm.

22. (B) For all x, P(x) is true.
The symbol representing the universal quantifier indicates that a given property or predicate applies to every single element within the specified domain of discourse. When prepended to a variable, it signifies that the statement following it must be true for all possible values that the variable can take. This is a foundational concept in predicate logic for expressing general rules and mathematical properties.

23. (D) Union
The union operation in set theory is logically analogous to the OR operator in Boolean logic. Just as the OR operator results in truth if at least one input is true, the union of two sets contains all elements that belong to the first set, the second set, or both. This relationship facilitates the translation between set-theoretic expressions and logical propositions.

24. (A) The basis step
The first essential phase in a proof by mathematical induction is the basis step. This involves verifying that the stated property or theorem holds true for the initial value in the sequence, which is typically the smallest natural number for which the statement is defined. Successfully completing this step provides the necessary foundation upon which the subsequent inductive argument is built.

25. (C) Every non-zero element has a multiplicative inverse.
A ring is an algebraic structure equipped with two binary operations, usually called addition and multiplication. While it requires an abelian group structure under addition and a semi-group structure under multiplication, it does not mandate the existence of multiplicative inverses for all non-zero elements. If a ring also includes multiplicative inverses for every non-zero element, it is classified more specifically as a field.

26. (A) Even
According to the handshaking lemma in graph theory, the sum of the degrees of all vertices in any graph is equal to twice the number of edges. Consequently, this sum must always be an even number. For this total to remain even, the number of vertices possessing an odd degree must be even, ensuring the overall parity of the graph is maintained.

27. (D) Disjunctive Normal Form (DNF)
Disjunctive Normal Form is a standard way of representing Boolean functions as a sum of one or more minterms. Each minterm represents a specific combination of input variables that results in a true output. This canonical representation is used extensively in digital logic design and automated theorem proving to provide a consistent and predictable structure for analyzing complex logical expressions and circuits.

28. (B) Time
The Program Evaluation and Review Technique is a project management tool primarily focused on managing the time aspect of complex projects. It utilizes probabilistic estimates for activity durations to account for uncertainty in project

Click Here to Buy book on Amazon

scheduling. By identifying the sequence of tasks that determine the minimum project completion time, PERT helps managers plan and monitor progress effectively to meet strict temporal deadlines.

29. (A) Both A and R are true and R is the correct explanation of A.
The assertion correctly identifies a logical equivalence where the conjunction operator is distributed over a disjunction. This specific structural transformation is defined by the distributive law in propositional logic, which mirrors the behavior of multiplication over addition in standard arithmetic. Since the reason accurately names and describes the principle behind the assertion, both are true and logically linked in this context.

30. (D) Reflexive, Anti-symmetric, and Transitive
A relation on a set qualifies as a partial ordering if it adheres to three specific criteria. It must be reflexive, meaning every element is related to itself. It must be anti-symmetric, ensuring that if two elements are mutually related, they are identical. Finally, it must be transitive, allowing the relationship to flow through sequences of elements. Together, these properties define posets.

31. (A) 0.7
To calculate the probability of the union of two events, one adds their individual probabilities and subtracts the probability of their intersection. This formula prevents double-counting the overlap between the events. In this case, adding zero point four and zero point five results in zero point nine, and subtracting the intersection of zero point two yields a final probability of zero point seven.

32. (B) $a * b^{-1}$ is in H
A non-empty subset of a group is considered a subgroup if it is closed under the group operation and contains the inverse of every element. A concise mathematical test for this condition states that for any two elements in the subset, the product of the first element and the inverse of the second element must also remain within the same subset for all cases.

33. (B) Dijkstra's algorithm
Dijkstra's algorithm is a widely used procedure in graph theory for determining the shortest path from a single source vertex to all other vertices in a weighted graph. It functions by iteratively selecting the unvisited vertex with the smallest tentative distance and updating the distances of its neighbors. This algorithm is particularly effective for graphs where all edge weights are non-negative values.

34. (A) XY + X'Z
The dual of a Boolean expression is created by interchanging the AND and OR operators and swapping the constant values zero and one. In the given expression, the addition operations are replaced with multiplications, and the multiplication between the parentheses is replaced with addition. This process does not involve complementing the variables themselves, but rather changes the structure of the logical operations.

35. (C) To minimize the total cost of shipping goods from sources to destinations.
The fundamental goal of a transportation model is to determine the most economical way to distribute goods from multiple sources to various destinations. It focuses on minimizing the total shipping cost while ensuring that the supply constraints of each source and the demand requirements of each destination are fully satisfied. This model is a specialized type of linear programming used in logistics.

36. (D) Modus Ponens
Modus Ponens is a fundamental rule of inference in propositional logic used to derive a conclusion from an implication and its antecedent. It states that if an if-then statement is accepted as true and the condition part of that statement is also true, then the consequence must follow. This logical step is essential for constructing valid deductive arguments in mathematical proofs.

37. (C) 64
The cardinality of a Cartesian product is the product of the cardinalities of the individual sets. With three elements in the first set and two in the second, the product contains six elements. The cardinality of the power set is calculated as two raised to the power of the number of elements

NODIA

in the original set. Thus, two to the sixth power is sixty-four.

38. (B) Finding the number of elements in the union of two or more sets.
The principle of inclusion-exclusion is a counting technique used to find the size of the union of multiple sets. It works by summing the sizes of individual sets, subtracting the sizes of all pairwise intersections, adding back the sizes of triple intersections, and continuing this alternating process. This method ensures that elements belonging to multiple sets are counted exactly once in the final total.

39. (D) It has no zero divisors.
An integral domain is a specific type of commutative ring with a unity element that possesses a critical property regarding zero. In such a structure, the product of any two non-zero elements can never result in zero. This absence of zero divisors allows for the cancellation of common non-zero factors in equations, similar to the behavior observed in the set of integers.

40. (A) It has no cycles of odd length.
A graph is classified as bipartite if its set of vertices can be divided into two disjoint groups such that no two vertices within the same group are adjacent. A fundamental theorem in graph theory states that a graph is bipartite if and only if it contains no cycles of odd length. This structural characteristic makes them useful for modeling relationships between different types.

41. (D) PERT Chart
While truth tables and Karnaugh maps are standard tools for representing and simplifying Boolean functions, and Venn diagrams can visually illustrate logical relationships between sets, a PERT chart is not used for this purpose. PERT charts are project management tools designed for scheduling and tracking tasks within a project timeline. They analyze the dependencies and durations of activities rather than logical bitwise operations.

42. (A) The effect of changes in input parameters on the optimal solution.
Sensitivity analysis in linear programming is a technique used to determine how changes in the input parameters of a model affect the optimal solution. This includes modifications to the coefficients of the objective function or the right-hand side values of the constraints. Understanding these impacts helps decision-makers evaluate the robustness of their results and prepare for potential fluctuations in real-world data.

43. (D) $p \rightarrow q$ is false only when p is true and q is false.
Logical connectives define how truth values are combined in propositions. A conjunction is true only if both parts are true, while a disjunction is false only if both parts are false. For a conditional implication, the statement is considered true in all scenarios except one. It specifically results in a false value only when the first part is true and the second is false.

44. (B) Reflexive
A relation defined on a set is called reflexive if every single element within that set bears the specified relationship to itself. This means that for any element in the domain, the pair consisting of that element and itself must be a member of the relation. Reflexivity is a core requirement for both equivalence relations and various types of ordered mathematical structures.

NODIA

45. (C) The probability of an event based on prior knowledge of conditions.
Bayes' theorem provides a mathematical framework for updating the probability of an event as new evidence or information becomes available. It calculates the conditional probability of a hypothesis based on prior knowledge of conditions that might be related to the event. This principle is fundamental in statistics, machine learning, and various fields where reasoning under uncertainty is required to make informed predictions.

46. (D) Endomorphism
In abstract algebra, a homomorphism is a map between two algebraic structures of the same type that preserves the operations. If the domain and the codomain of this mapping are the exact same group, the homomorphism is specifically referred to as an endomorphism. This type of mapping explores the internal structural relationships and transformations that can occur within a single algebraic object itself.

47. (A) Contains all the vertices of G and is a tree.
A spanning tree of a connected graph is a subgraph that includes every vertex of the original graph while maintaining a tree structure. This means the subgraph must be connected and contain no cycles. Spanning trees are critical in network design for finding the most efficient way to connect all nodes with the minimum number of edges, often used in routing protocols and optimization.

48. (B) AB'C
A minterm is a product term that evaluates to true for exactly one specific combination of input variables. To represent the binary input one-zero-one for variables A, B, and C, each variable is included in its true form if the bit is one and in its complemented form if the bit is zero. Therefore, the resulting term consists of A, B-prime, and C.

49. (B) Hungarian Method
The Hungarian method is a combinatorial optimization algorithm specifically designed to solve assignment problems in polynomial time. It works by manipulating the cost matrix through row and column reductions to identify a set of zeros that represent an optimal assignment. This technique is highly efficient for matching tasks to resources in a way that minimizes the total cost or maximizes the total effectiveness.

50. (A) P-4, Q-2, R-1, S-3
In propositional logic, specific names are assigned to fundamental equivalences. The complement law involves a variable and its negation. The commutative law states that the order of operands does not change the result. De Morgan's laws describe how to negate conjunctions and disjunctions. Finally, the absorption law demonstrates how a variable can absorb a redundant term, simplifying the overall logical expression significantly.

51. (C) A poset
A set equipped with a partial ordering relation is formally known as a partially ordered set, or more commonly, a poset. This structure allows for some elements to be compared with others while leaving some pairs potentially incomparable. Posets are fundamental in mathematics for representing hierarchies, schedules, and structural relationships where a consistent but not necessarily complete order exists between the various elements.

52. (B) Both Statement I and Statement II are false.
Permutations and combinations are distinct counting methods in combinatorics. A permutation refers to an arrangement where the specific order of objects is vital. In contrast, a combination refers to a selection of objects where the order of choosing does not matter. Since both provided statements swapped these definitions, they are both incorrect descriptions of how these mathematical concepts are applied in counting scenarios.

53. (C) An automorphism
An automorphism is a special type of isomorphism where the mapping starts and ends within the same group. It represents a way to rearrange the elements of the group such that the underlying algebraic structure remains perfectly preserved. Studying automorphisms helps mathematicians understand the symmetries inherent in a group and how its elements can be transformed while maintaining the original group operation's properties.

54. (B) Hamiltonian circuit problem
The Hamiltonian circuit problem involves finding a path in a graph that starts and ends at the same vertex and visits every other vertex exactly once. This is different from an Eulerian circuit, which focuses on edges. Determining the existence of such a circuit is a classic problem in computer science and graph theory, often associated with optimization and the traveling salesperson problem.

55. (D) NAND gate
A universal logic gate is one that can be used to implement any other Boolean function without the need for additional gate types. The NAND gate is a prime example of such a gate. By combining multiple NAND gates in specific configurations, one can recreate the behavior of AND, OR, and NOT gates, which are sufficient to build any complex digital circuit.

56. (A) The initial simplex tableau is optimal but infeasible.
The dual simplex method is an alternative approach to the standard simplex algorithm used

Downloaded from www.pdf.tube

in linear programming. It is specifically applied when the current solution is optimal, meaning the optimality condition is satisfied, but it is currently infeasible because some basic variables have negative values. This method works by iteratively moving through infeasible solutions toward feasibility while maintaining the optimal condition throughout.

57. (C) Inference
Inference is the logical process used to derive new true statements or conclusions from a set of existing facts or premises. It involves applying valid rules, such as Modus Ponens or Modus Tollens, to progress through a series of logical steps. This mechanism is the core of deductive reasoning and is used to construct formal proofs and solve problems in mathematics.

58. (A) The order of elements in a set is important.
In the mathematical definition of a set, the collection of elements is unordered. This means that listing the same elements in a different sequence does not create a new or different set. Other properties, such as the empty set being a universal subset and the power set cardinality being two raised to the number of elements, are correct and fundamental to set theory.

59. (D) Their intersection is the empty set.
In probability theory, two events are considered mutually exclusive if it is impossible for both to occur at the same time. Mathematically, this means that the intersection of the two event sets is an empty set, resulting in a joint probability of zero. This concept is vital for correctly applying the addition rule when calculating the probability that either event will happen.

60. (B) Commutative ring
A ring is a set with addition and multiplication operations. While addition in a ring must be commutative, multiplication does not necessarily have this property. If a ring is structured such that the order of factors in a multiplication does not change the product for any two elements, it is specifically called a commutative ring. This property is common in many familiar number systems.

61. (C) The minimum number of colors needed to color the vertices such that no two adjacent vertices share the same color.
The chromatic number of a graph represents the fewest number of colors required to color each vertex such that no two vertices connected by an edge share the same color. This concept is a central problem in graph theory with applications in scheduling, register allocation, and map coloring. Finding the chromatic number is generally computationally difficult for large and complex graphs.

62. (B) Conjunctive Normal Form (CNF)
Conjunctive Normal Form is a standard representation of Boolean functions where the expression is written as a product of one or more sum terms. This structure is also known as Product of Sums. Each sum term, or clause, consists of variables or their negations combined with the OR operator. This form is particularly useful for certain types of logical proofs and circuit designs.

63. (A) Some or all of the decision variables are restricted to be integer values.
Integer programming is a specialized branch of mathematical optimization where the decision variables are required to take on only integer values. This constraint is added to a standard linear programming model to reflect situations where fractional solutions are not physically or logically possible, such as when counting items or people. These problems are often significantly more difficult to solve than continuous ones.

NODIA

64. (B) (i), (iii), (iv)
For a set and a binary operation to form a group, four axioms must be satisfied: closure, associativity, the existence of an identity element, and the existence of an inverse for every element. While some groups are also commutative, this property is not a requirement for the basic definition of a group. Therefore, identity, inverses, and associativity are the essential criteria for group membership.

65. (B) Its variables are bound by quantifiers or assigned specific values.
A predicate is a statement containing variables that is neither true nor false until those variables are specified. To transform a predicate into a

proposition, which has a definite truth value, the variables must be either assigned specific values from the domain or bound by quantifiers such as for all or there exists. This process removes the ambiguity inherent in the variable.

66. (A) (a, b) ∈ R and (b, a) ∈ R implies a = b for all a, b ∈ A.
A relation on a set is anti-symmetric if the only way two elements can be related to each other in both directions is if they are actually the same element. This property ensures that no distinct pairs of elements exhibit a reciprocal relationship. Anti-symmetry is a defining characteristic of partial orderings, such as the less than or equal to relation on numbers.

67. (C) n
The notation C(n, 1) represents the number of ways to choose exactly one object from a set of n distinct objects. According to the formula for combinations, this is calculated by dividing n factorial by the product of one factorial and (n-1) factorial. This simplification always results in n, as there are exactly n individual choices available when selecting a single item.

68. (D) The longest path in the network, which determines the minimum project completion time.
In project management techniques like PERT and CPM, the critical path is the longest sequence of dependent activities through a network. The total duration of this path determines the minimum time required to complete the entire project. Any delay in an activity on this path will directly result in a delay of the project completion, making these tasks the most time-critical for managers.

69. (C) Planar graph
A planar graph is a type of graph that can be embedded in a two-dimensional plane such that its edges intersect only at their endpoints. This means no two edges cross over each other when the graph is drawn. Determining whether a graph is planar is important in fields like circuit board design and cartography, where avoiding edge crossings is often a requirement.

70. (A) Both A and R are true and R is the correct explanation of A.
In group theory, the identity element is defined as the element that leaves any other element unchanged under the group operation. By applying this definition to the identity itself, it follows that the identity operated with itself must yield the identity. Since an inverse is defined as an element that results in the identity, the identity element must be its own unique inverse.

71. (D) A+B'+C
A maxterm is a sum term that evaluates to false for exactly one specific combination of input variables. To represent the binary input zero-one-zero for variables A, B, and C, each variable is included in its true form if the bit is zero and in its complemented form if the bit is one. Therefore, the resulting term consists of A plus B-prime plus C.

72. (A) P(k) is true for some arbitrary integer k.
During the inductive step of a proof by mathematical induction, the statement is assumed to be true for an arbitrary but fixed integer, typically denoted as k. This assumption is known as the inductive hypothesis. The goal of the inductive step is then to use this hypothesis to prove that the statement must also be true for the next integer in the sequence.

73. (C) Implication Equivalence
The logical equivalence between an implication and the disjunction of the negated antecedent and the consequent is a fundamental rule in propositional logic. This identity, often referred to as implication equivalence, allows complex logical expressions to be rewritten using only negations and disjunctions. This transformation is particularly useful when converting formulas into standard forms like CNF or DNF for computational analysis.

74. (A) Even degree
A connected graph possesses an Eulerian circuit, which visits every edge exactly once and returns to the start, if and only if every vertex in the graph has an even degree. This requirement ensures that whenever a path enters a vertex via one edge, there is always an available unused edge to exit that same vertex, eventually allowing the traversal to close.

Click Here to Buy book on Amazon

75. (D) To visit every node in the tree in a systematic way.
Tree traversal algorithms, such as in-order, pre-order, and post-order, are designed to visit every node within a tree structure in a specific, predictable order. These systematic methods are fundamental for various tasks in computer science, including searching for data, printing the contents of the tree, and performing evaluations on expression trees, ensuring that no node is missed or processed incorrectly.

76. (B) It is a commutative ring in which every non-zero element has a multiplicative inverse.
A field is a sophisticated algebraic structure that builds upon the properties of a ring. It is defined as a commutative ring with unity where every single non-zero element has a corresponding multiplicative inverse. This property allows for the operations of addition, subtraction, multiplication, and division to be performed consistently, similar to the behavior of rational, real, or complex number systems.

77. (C) Transportation problem
The assignment problem is a specialized version of the transportation problem where the goal is to allocate exactly one resource to exactly one task. In this model, the supply at each source and the demand at each destination are all set to one. Consequently, any algorithm used to solve transportation problems can also be applied to find the optimal solution for assignments.

78. (D) A contradiction
A contradiction is a compound proposition that is always false, regardless of the truth values assigned to its individual components. This occurs when the logical structure of the statement inherently conflicts with itself, such as in the case of a variable being both true and false simultaneously. Contradictions are used in proofs by contradiction to show that an initial assumption was incorrect.

79. (A) The power set of A
The power set of any given set is the collection of all possible subsets that can be formed from the elements of that set, including the set itself and the empty set. If the original set contains n elements, the power set will contain two raised to the power of n elements. Power sets are fundamental in set theory and for defining various structures.

80. (B) n! / (n-r)!
The number of ways to arrange r objects selected from a total of n distinct objects is calculated using the permutation formula. This value is determined by taking the factorial of n and dividing it by the factorial of the difference between n and r. This accounts for the specific order of the objects in the arrangement, which is a key requirement for permutations.

NODIA

81. (B) (ii), (iii), (v) only
In any mathematical group, the identity element and the inverse of each individual element are proven to be unique. Additionally, the cancellation laws are valid for all elements within the structure. While integers under addition satisfy the group axioms, natural numbers do not because they lack additive inverses. Furthermore, by definition, a group must have exactly one unique identity element, never more.

82. (A) A leaf
In the terminology of rooted trees, a leaf is defined as any vertex that has no children. These nodes represent the terminal points of the tree branches. In contrast, nodes with at least one child are known as internal nodes, while the unique top-level node with no parent is called the root. Identifying leaves is essential for many recursive algorithms and data structures.

83. (C) Exactly one variable.
The effectiveness of a Karnaugh map for simplifying Boolean expressions is derived from the adjacency of its cells. The map is organized so that moving between adjacent cells corresponds to a change in exactly one variable. By grouping these adjacent cells containing ones, variables that change while the output remains the same can be eliminated, resulting in a significantly reduced and simplified expression.

84. (D) Zero or negative
Slack, also known as float, represents the amount of time an activity can be delayed without affecting the overall project schedule. For activities located on the critical path, there is no flexibility in timing. Consequently, the slack for these specific activities is always zero. In some

project scenarios with fixed deadlines, this value can also be negative, indicating the project is currently behind.

85. (C) $\forall x \exists y P(x, y)$
Nested quantifiers occur in predicate logic when one quantifier is placed within the scope of another quantifier. This structure allows for the expression of complex relationships involving multiple variables, such as for every x, there exists a y. The order and type of these quantifiers significantly change the meaning of the statement, which is essential for formalizing various mathematical definitions and properties.

86. (B) Anti-symmetric
An equivalence relation is defined by the three properties of reflexivity, symmetry, and transitivity. Anti-symmetry is a property where if two elements relate to each other in both directions, they must be the same element. This property is characteristic of partial orders and is generally incompatible with symmetry unless the relation only includes self-loops. Therefore, anti-symmetry is not a requirement for equivalence.

87. (D) 10
To find the number of ways to choose three books from a set of five without regard to order, the combination formula is used. This is calculated by dividing five factorial by the product of three factorial and two factorial. Multiplying out the values and performing the division results in ten unique combinations of books that can be selected from the given collection.

88. (A) A set of edges whose removal increases the number of connected components.
In graph theory, a cut-set is a collection of edges whose removal from a connected graph results in an increase in the number of its connected components. Effectively, removing these edges disconnects the graph into two or more separate pieces. This concept is fundamental for analyzing the reliability, connectivity, and maximum flow capacity of networks, such as communication or transportation systems.

89. (C) Prefix code
A prefix code is a type of variable-length code where no valid codeword in the set is a prefix of any other codeword. This property ensures that the code is uniquely decodable without needing a special separator between symbols. These codes are also referred to as instantaneous codes because each symbol can be decoded as soon as its last bit is received.

90. (A) All constraints, including the non-negativity restrictions.
The feasible region in a linear programming problem is the set of all possible points that satisfy every constraint simultaneously. This includes all the linear inequalities and equalities defined in the model, as well as the non-negativity restrictions placed on the decision variables. Any point within this region represents a potential solution, and the optimal solution is typically found at its boundaries.

91. (A) P-2, Q-3, R-1, S-4
Algebraic structures are classified by the axioms they satisfy. A semi-group requires only closure and associativity. A monoid adds the requirement of an identity element. A group further requires that every element has an inverse. Finally, an Abelian group must satisfy all these properties plus commutativity. Matching these structures to their specific requirements provides a hierarchical understanding of various algebraic systems.

92. (B) $\wedge$
In formal logic, the conjunction of two propositions is represented by the wedge symbol, which resembles an inverted letter 'v'. This operator corresponds to the word and in natural language. A conjunction results in a true value only if both of the connected propositions are true simultaneously. It is a fundamental building block for creating more complex compound statements in propositional logic.

93. (C) Reflexive, Anti-symmetric, and Transitive
The divides relation on positive integers is reflexive because every number divides itself. It is anti-symmetric because if a divides b and b divides a, then a and b must be equal. It is also transitive, meaning if a divides b and b divides c, then a divides c. These three properties define the relation as a partial ordering on the set.

94. (D) 1/13

Downloaded from www.pdf.tube

A standard deck of playing cards contains fifty-two cards in total, divided into four suits. Each suit contains exactly one ace, resulting in a total of four aces in the entire deck. The probability of drawing an ace is calculated by dividing the number of favorable outcomes, four, by the total number of possible outcomes, fifty-two. This simplifies to one thirteenth.

95. (A) The sum of the weights of its constituent edges is minimized.
In the context of a weighted graph, the shortest path problem seeks to find a sequence of edges between two nodes such that the total sum of the weights associated with those edges is minimized. The weight can represent various metrics such as distance, time, or cost. This is a fundamental optimization problem with numerous applications in navigation, network routing, and logistics planning.

96. (B) Canonical Sum of Products (Disjunctive Normal Form)
The Canonical Sum of Products, also known as the Disjunctive Normal Form, provides a unique and standardized representation for any Boolean function. It expresses the function as a sum of minterms, where each minterm corresponds to a specific row in the truth table where the function evaluates to true. This representation is essential for comparing functions and for certain types of circuit implementation.

97. (B) Adjusting the schedule to reduce peaks and troughs in resource demand.
Resource levelling is a project management technique used to address resource conflicts and over-allocations. It involves adjusting the project schedule by shifting tasks within their available slack time to ensure that the demand for resources remains steady and does not exceed availability. This process helps create a more balanced workload and can prevent the need for costly resource acquisitions or overtime.

98. (A) Has a specific syntactic structure, like CNF or DNF.
In propositional logic, a normal form is a standardized syntactic structure for logical formulas. Two of the most common types are Conjunctive Normal Form and Disjunctive Normal Form. Converting a formula into a normal form involves using logical equivalences to ensure it follows a specific pattern of operators. This standardization is crucial for automated reasoning and simplifying the process of checking for tautologies.

NODIA

99. (C) Division
Standard set theory defines several basic operations to combine or modify sets, including union, intersection, and complement. Union combines all elements, intersection finds shared elements, and complement identifies elements outside a set. However, division is not a standard set operation. While specific contexts like group theory might involve division-like concepts, it does not exist as a primary operation in basic set algebra.

100. (D) Inverse
In the context of a group with a binary operation and an identity element, if the combination of two elements results in the identity, those elements are considered inverses of each other. Every element in a group is required to have a unique inverse. Applying an element followed by its inverse essentially reverses the first operation, returning the system to its original state.

101. (C) n(n-1)/2
A complete graph with n vertices is one where every pair of distinct vertices is connected by a unique edge. To find the total number of edges, one calculates the number of ways to choose two vertices from the set of n vertices. This is given by the combination formula n times (n minus one) divided by two, representing all possible pairwise connections.

102. (A) Both Statement I and Statement II are true.
Both statements regarding tree traversals are correct definitions of how nodes are visited. Post-order traversal specifically processes all child subtrees before visiting the parent node, ensuring the root is the last node processed. Conversely, pre-order traversal visits the parent node before moving on to its children, making the root the first node visited. These patterns are fundamental to various tree algorithms.

103. (D) Ensuring that resource usage is kept at a constant level during the project.

Resource levelling in project management techniques like PERT and CPM is the process of smoothing out resource requirements over time. Its primary goal is to ensure that the demand for specific resources, such as personnel or equipment, stays as constant as possible throughout the project duration. This avoids large fluctuations in resource needs, making management more predictable and reducing the risk of shortages.

104. (A) There exists an integer x such that $x \leq 5$.
To negate a universally quantified statement, the for all quantifier is replaced with there exists, and the property within the statement is negated. The negation of being greater than five is being less than or equal to five. Therefore, the negation of the claim that all integers are greater than five is that there exists at least one integer less than or equal to five.

105. (C) Disjoint equivalence classes
An equivalence relation defined on a set naturally divides that set into distinct, non-overlapping subsets known as equivalence classes. Every element of the original set belongs to exactly one class. All elements within a specific class are related to each other by the equivalence relation, while no elements from different classes are related. This partitioning is a fundamental concept in abstract mathematics.

106. (A) 35
Forming a committee from a group of people is a combination problem because the order in which the members are chosen does not change the resulting committee. To find the number of ways to pick three people from seven, one divides seven factorial by the product of three factorial and four factorial. Calculating these values yields thirty-five different possible committees that can be formed.

107. (D) Every group has at least two subgroups: the group itself and the trivial subgroup.
Every group possesses at least two subgroups, provided it is not the trivial group containing only the identity. These are the group itself and the trivial subgroup consisting solely of the identity element. Other statements regarding the union or intersection of subgroups are not universally true. For example, the union of two subgroups is only a subgroup if one is contained within the other.

108. (B) In-order
A Binary Search Tree is organized such that the value of any node is greater than the values in its left subtree and smaller than the values in its right subtree. When an in-order traversal is performed, the algorithm visits the left child, then the node, then the right child. This specific sequence results in the nodes being processed in non-decreasing, sorted order.

109. (C) Hamming weight
The Hamming weight of a binary string or a number is the count of symbols that are different from the zero symbol. In the context of binary representation, this corresponds exactly to the number of bits that are set to one. This metric is widely used in information theory, coding theory, and cryptography to measure the density of information and detect errors.

110. (D) Exactly two vertices have odd degree.
A connected graph contains an Eulerian path if and only if it has exactly zero or two vertices with an odd degree. If the graph has exactly two odd-degree vertices, the path must start at one of these vertices and end at the other. If the graph has no odd-degree vertices, it contains an Eulerian circuit, which is a special type of path.

111. (C) All (i), (ii), (iii), (iv), and (v) are correct.
Boolean algebra is governed by a specific set of identities and laws that define how binary variables interact. Properties such as identity elements for addition and multiplication, complementation resulting in zero or one, and idempotency where a variable added to itself remains unchanged are all fundamental. Each of the provided statements accurately reflects these core principles, which are essential for simplifying digital logic.

112. (B) Decision variables, Objective function, Constraints
A mathematical model for a linear programming problem is composed of three essential parts. First, decision variables represent the quantities to be determined. Second, the objective function defines the goal of maximizing or minimizing a linear expression of these variables. Third,

Click Here to Buy book on Amazon

a set of linear constraints and non-negativity restrictions delineate the limits on resources and the boundaries of the feasible solution space.

113. (A) 1, 3, 4, 2, 5
The Simplex method follows a structured iterative process. It begins by constructing the initial tableau. The algorithm then checks for optimality to see if improvement is possible. If not, it identifies which variable should enter and leave the basis to improve the objective value. After selecting the pivot element, row operations are performed to update the tableau before repeating the optimality check.

114. (A) Contingency
In logic, propositions are classified based on their truth tables. A tautology is always true, while a contradiction is always false. A contingency refers to a proposition that is true in some scenarios and false in others, depending on the truth values of its variables. Most everyday statements are contingencies, as their truth depends on the specific context of the underlying variables.

115. (C) R is not symmetric but is transitive.
The relation is a sister of is not symmetric because if Alice is the sister of Bob, Bob is not the sister of Alice; he is her brother. However, the relation is transitive because if Alice is the sister of Brenda, and Brenda is the sister of Chloe, then Alice must also be the sister of Chloe, assuming they share the same parents.

116. (D) Mathematical induction
Mathematical induction is a powerful proof technique used to establish the truth of a statement for all natural numbers. It consists of two steps: the basis step, which proves the statement for the first number, and the inductive step, which shows that if the statement holds for one number, it must hold for the next. Together, these steps create a chain of proof.

117. (C) $ghg^{-1} \in H$
A subgroup is defined as normal if it is invariant under the conjugation operation by any element of the parent group. Mathematically, this means that for every element in the subgroup and every element in the group, the result of the group element times the subgroup element times the inverse of the group element must remain within the subgroup. This property allows for well-defined quotient groups.

118. (B) has no cycles
In graph theory, a tree is characterized by being a connected graph that contains no cycles. This acyclic nature means there is exactly one unique path between any two vertices in the graph. Trees are a fundamental type of graph used to represent hierarchical relationships and are the most efficient way to connect a set of vertices with the minimum number of edges.

119. (D) All of the above
Boolean functions can be effectively simplified by applying a variety of algebraic laws. Associative and commutative laws allow for the regrouping and reordering of terms. The distributive law enables the expansion or factoring of expressions. Idempotent and absorption laws further reduce redundancy by eliminating repeated variables or terms that are logically contained within others. Using all these laws together optimizes logical expressions.

120. (A) (to + 4tm + tp) / 6
The Program Evaluation and Review Technique calculates the expected duration of an activity using a weighted average of three time estimates. This formula gives the most likely time four times the weight of the optimistic and pessimistic estimates, and then divides the total by six. This approach provides a more balanced and realistic time estimate for activities that involve a degree of uncertainty.

121. (C) Modus Tollens
Modus Tollens is a valid rule of inference that allows one to conclude the negation of an antecedent if the consequent of an implication is known to be false. If if p then q is true, and not q is also true, then it logically follows that not p must be true. This rule is a core component of deductive reasoning in mathematics.

122. (A) $\{x \mid x \in A \text{ and } x \notin B\}$
The difference between two sets, A and B, is a set containing all elements that are members of A but are not members of B. This operation essentially removes the overlap between the two sets from the first set. It is a fundamental set operation used to isolate unique elements and is often used

Downloaded from www.pdf.tube

in conjunction with unions and intersections to define complex sets.

123. (A) Both A and R are true and R is the correct explanation of A.
The generalized version of the Pigeonhole Principle states that if a large number of items are placed into a smaller number of containers, at least one container must hold multiple items. The specific case where the number of pigeons exceeds the number of holes is the most basic application of this rule. The reasoning provided correctly identifies the underlying principle that governs these distributions.

124. (B) The set of rational numbers under addition and multiplication.
A field requires that every non-zero element has a multiplicative inverse within the same set. The set of rational numbers satisfies this because for any non-zero fraction, its reciprocal is also a rational number. In contrast, integers do not form a field because the multiplicative inverse of an integer like two is not an integer. Rational numbers also satisfy all other field axioms.

125. (C) Kruskal's algorithm
Kruskal's algorithm is a greedy method used to find the minimum spanning tree of a connected, weighted graph. It works by sorting all edges by their weights and adding them to the tree one by one, provided that adding an edge does not create a cycle. This process continues until all vertices are connected, resulting in a tree with the minimum total edge weight.

126. (D) Sum of Products (SOP) form
A Boolean function represented as a collection of terms joined by addition, where each term is a product of variables or their negations, is in Sum of Products form. In the given function, several minterms are summed together. This format is a standard way to express logical functions and is directly related to the rows of a truth table where the output is true.

127. (A) PERT uses a probabilistic time estimate while CPM uses a deterministic time estimate.
The fundamental distinction between PERT and CPM lies in how they handle activity durations. PERT is a probabilistic technique that uses three different time estimates to account for uncertainty in task completion. In contrast, CPM is a deterministic method that assumes task durations are known with certainty. This difference makes PERT better suited for research projects and CPM better for repetitive construction tasks.

128. (B) $(p \rightarrow q) \wedge (q \rightarrow p)$
A biconditional proposition, p if and only if q, is logically equivalent to the conjunction of two conditional statements: if p then q, and if q then p. This means that for the biconditional to be true, the implication must hold in both directions simultaneously. This equivalence is often used to break down complex biconditional proofs into two simpler, directional implication proofs.

129. (B) {(a, b), (c, d)}
A relation on a set A is defined as a subset of the Cartesian product of A with itself. Since the set A only contains elements a, b, and c, any valid relation on A can only consist of ordered pairs containing these three elements. A relation containing the element d is invalid because d is not a member of the original set A provided.

130. (A) The statement P(k) is true for an arbitrary but fixed integer k.
The inductive hypothesis is the core assumption made during the second step of a proof by mathematical induction. It involves assuming that the statement to be proven is true for a specific, but arbitrarily chosen, integer k. This assumption provides the necessary starting point to demonstrate that the truth of the statement for k logically necessitates its truth for the subsequent integer k plus one.

131. (B) P-2, Q-1, R-4, S-3
Different graph algorithms rely on specific data structures for their operation. Breadth-First Search uses a queue to explore nodes level by level. Depth-First Search utilizes a stack to go as deep as possible along branches. Shortest path algorithms like Dijkstra's and minimum spanning tree algorithms like Prim's both use priority queues, but they prioritize different metrics to determine which node to process next.

132. (D) A commutative ring with unity and no zero-divisors.

An integral domain is a specific type of commutative ring that possesses a multiplicative identity and lacks zero divisors. This means that if the product of two elements is zero, then at least one of those elements must itself be zero. This property is crucial because it allows for the cancellation of non-zero factors in algebraic equations, making the structure behave similarly to the familiar set of integers.

133. (D) Statement I is false, but Statement II is true.
A simple graph is restricted to having at most one edge between any two vertices and no edges connecting a vertex to itself. In contrast, a multigraph allows for multiple edges between the same pair of vertices, but it does not necessarily include loops unless it is specifically defined as a pseudograph. However, in many contexts, the definition of a multigraph is broadened to include loops.

134. (B) Time and cost
In project scheduling, cost consideration focuses on the relationship between the time taken to complete a project and the total cost incurred. Crashing is a common technique where additional resources are allocated to critical activities to shorten their duration, which usually increases the direct costs. Managers must analyze this trade-off to find the most cost-effective way to meet a project deadline or schedule.

135. (D) $\forall x\,(M(x) \rightarrow A(x))$
To translate the statement All mammals are animals into predicate logic, a universal quantifier is used to cover all individuals in the domain. The statement is structured as an implication, asserting that if any given individual is a mammal, then that same individual must also be an animal. This is represented by a universal quantifier followed by the implication from the mammal to animal predicate.

136. (A) A partial ordering
The is a subset of relation on a power set is a classic example of a partial ordering. It is reflexive because every set is a subset of itself. it is anti-symmetric because if two sets are subsets of each other, they are identical. Finally, it is transitive because if A is a subset of B and B of C, A is of C.

137. (C) m * n
The fundamental principle of counting, also known as the product rule, states that if one event can occur in m ways and a second independent event can occur in n ways, then the total number of ways the two events can occur in sequence is the product of m and n. This rule is the basis for calculating permutations and other complex counting problems.

138. (A) An algebraic structure formed by partitioning a larger structure by an equivalence relation.
A quotient of an algebraic structure is a new structure formed by grouping the original elements into equivalence classes based on a specific congruence relation. This process essentially divides the structure by an equivalence relation that is compatible with the underlying operations. For example, quotient groups are formed by partitioning a group into cosets of a normal subgroup, preserving the group's operational properties.

139. (D) The vertices can be partitioned into two disjoint sets, V1 and V2, such that every edge connects a vertex in V1 to one in V2.
A graph is defined as bipartite if its vertex set can be split into two separate, non-overlapping groups. The critical requirement for this structure is that every edge in the graph must connect a vertex from the first group to a vertex from the second group. There can be no edges between vertices within the same group. This makes bipartite graphs ideal for representing relationships between sets.

140. (B) $(x + y)' = x' + y'$
De Morgan's Laws describe how to negate conjunctions and disjunctions. The correct law states that the negation of a sum is equal to the product of the negations. Therefore, the complement of the sum of x and y should be the product of the complement of x and the complement of y. The provided statement incorrectly suggests that the negation of a sum remains a sum.

141. (C) Assignment problems
The Hungarian method is a well-known combinatorial optimization algorithm specifically designed to solve the assignment problem. Its primary function is to find the minimum cost or maximum profit way to assign a set of tasks

Click Here to Buy book on Amazon

to a set of agents, where each task is assigned to exactly one agent. This method is highly efficient and operates within polynomial time for most practical applications.

142. (B) (i), (iii), (v)
Several core theorems define graph theory. A simple graph with n vertices has a maximum of n times (n minus one) divided by two edges. The handshaking lemma confirms the sum of degrees is twice the edge count. Furthermore, every connected graph contains at least one spanning tree. However, trees with n vertices must have exactly n minus one edges, and not all complete graphs are planar.

143. (A) $\neg\forall x\ P(x) \equiv \exists x\ \neg P(x)$
De Morgan's Laws for quantifiers provide a method for negating quantified statements in predicate logic. The negation of a for all statement is logically equivalent to saying there exists at least one element for which the property does not hold. Similarly, negating an existential statement is equivalent to claiming the property is false for all elements. These equivalences are vital for simplifying complex logical expressions.

144. (B) Smallest relation containing R that is transitive.
The transitive closure of a relation is the smallest possible relation that contains all the pairs of the original relation and also satisfies the property of transitivity. To form the transitive closure, one must add all the indirect connections implied by the existing pairs. This resulting relation is the most minimal extension that ensures if a relates to b and b to c, a relates to c.

145. (B) Cardinality of the union of sets.
The inclusion-exclusion principle is a fundamental counting technique used to determine the total number of elements in the union of several finite sets. It avoids overcounting elements that belong to more than one set by systematically adding the sizes of individual sets and then subtracting and adding the sizes of their various intersections. This principle is widely used in probability and combinatorics problems.

146. (A) Monoid
A monoid is an algebraic structure that sits between a semi-group and a group in the hierarchy of complexity. It consists of a set and a binary operation that is both closed and associative. The defining characteristic that distinguishes a monoid from a basic semi-group is the additional requirement that the set must contain a unique identity element for the given operation.

147. (C) n-1
A fundamental property of any tree with n vertices is that it must contain exactly n minus one edges. This specific edge count is the minimum required to connect all n vertices while simultaneously ensuring that no cycles are formed. If any edge is added, a cycle will be created; if any edge is removed, the graph will become disconnected and no longer be a tree.

148. (D) Interchanging 0s and 1s, and interchanging AND and OR operators.
To find the dual of a Boolean expression, one must follow two specific replacement rules. Every AND operator is replaced with an OR operator, and every OR operator is replaced with an AND operator. Additionally, every occurrence of the constant zero is swapped with the constant one, and vice versa. Variables and their complements are left unchanged during this specific structural transformation of the expression.

NODIA

149. (C) Both the rim conditions and non-negativity constraints
In a transportation problem, a solution is considered feasible if it satisfies all the constraints of the model. This includes the rim conditions, which ensure that the total amount shipped from each source equals its supply and the total received at each destination equals its demand. Additionally, all shipping quantities must satisfy non-negativity constraints, meaning no negative amounts of goods can be transported.

150. (A) 2, 3, 1, 4
Different algebraic structures require varying numbers of axioms. A field is the most complex, requiring addition and multiplication groups. A ring follows, requiring an additive group and a multiplicative semi-group with distribution. A group requires four axioms: closure, associativity, identity, and inverse. Finally, a semi-group is the simplest, requiring only closure and associativity. This ranking reflects the increasing mathematical restrictions placed on each structure.

151. (D) Existential Instantiation
Rules of inference like Modus Ponens, Resolution, and Conjunction Introduction are used in propositional logic to derive new true statements from existing ones. Existential Instantiation is a rule used in predicate logic, not propositional logic, as it deals with quantifiers. This rule allows one to conclude that a specific element exists if an existential statement is known to be true for the domain.

152. (A) An equivalence relation on a set causes the set to be partitioned into disjoint equivalence classes.
There is a direct mathematical relationship between equivalence relations and set partitions. An equivalence relation on a set is the mechanism that causes the set to be divided into a collection of disjoint equivalence classes. These classes together form a partition of the original set, where every element is accounted for and no element belongs to more than one class. This is a fundamental concept.

153. (C) 4/10
The probability of an event is calculated by dividing the number of favorable outcomes by the total number of possible outcomes. In this scenario, there are four red balls and ten total balls in the bag. Therefore, the probability of selecting a red ball is four divided by ten. This fraction represents the likelihood of drawing a red ball in a single random selection.

154. (A) 0
An identity element for a binary operation is an element that leaves any other element unchanged when combined with it. For the set of integers under the operation of addition, zero serves as the identity because adding zero to any integer results in the same integer. This property is one of the requirements for integers to form a group under the addition operation.

155. (D) Assigning colors to vertices such that adjacent vertices have different colors.
Graph coloring is a method of assigning labels, traditionally called colors, to the elements of a graph subject to certain constraints. In its most common form, vertex coloring, the goal is to color each vertex such that no two adjacent vertices share the same color. The minimum number of colors required for this task is a fundamental property of the graph known as its chromatic number.

156. (B) Two variables eliminated
In a Karnaugh map, grouping adjacent cells containing ones allows for the simplification of Boolean expressions. Each time the size of the group doubles, another variable is eliminated from the resulting term. A group of two adjacent ones eliminates one variable, while a group of four adjacent ones eliminates two variables. This occurs because those two variables change values within the group without affecting the output.

157. (C) The optimal solution to the relaxed LP problem provides a bound for the integer solution.
Solving an integer programming problem is generally more complex than solving its corresponding linear programming relaxation, where variables are allowed to be fractional. However, the optimal solution to the relaxed LP problem is useful because it provides a mathematical bound for the objective value of the true integer solution. This bound is essential for algorithms like branch and bound to efficiently search the solution space.

158. (D) Tautology
A tautology is a statement that is true for every possible combination of truth values of its variables. In this specific implication, if the conjunction part is true, then both p and q must be true, which automatically makes the disjunction part true. Since it is impossible for the first part to be true while the second is false, the entire statement is a tautology.

159. (A) Cardinality
In the context of database systems and set theory, a relation is defined as a set of tuples. The cardinality of a relation refers specifically to the total number of these tuples within the set. This is a measure of the size of the relation in terms of the number of individual records or data instances it contains, distinguishing it from the degree or arity.

160. (A) Both A and R are true and R is the correct explanation of A.
The formula for the number of permutations of n distinct objects is indeed n factorial. This

Downloaded from www.pdf.tube

is logically derived using the product rule of counting, as there are n choices for the first position, n minus one for the second, and so on until one choice remains. Since the reasoning accurately explains the combinatorial origin of the assertion, both are true and properly linked.

161. (B) (A, *) satisfies the closure property
A non-empty set combined with a binary operation is considered an algebraic structure if it satisfies the closure property. This property ensures that when the operation is applied to any two elements from the set, the resulting value is also an element of the same set. Closure is the most basic requirement for defining any formal system where operations can be repeatedly and consistently applied.

162. (A) Weighted Graph
A weighted graph is a type of graph where each edge is assigned a numerical value, known as a weight or cost. These weights can represent physical distances, traversal times, financial costs, or any other relevant metric. Weighted graphs are essential for modeling and solving optimization problems, such as finding the shortest path between two locations or the minimum cost to connect a network.

163. (C) x'
A Boolean function F(x, y, z) defined by the sum of minterms zero through three indicates that the function is true whenever the variable x is zero, regardless of the values of y and z. In a truth table, these four minterms correspond to all combinations where the first bit is zero. Therefore, the function simplifies to the complement of the first variable, x-prime.

164. (D) To start with an infeasible but better than optimal solution and move towards feasibility.
The dual simplex method is designed to handle linear programming problems where the initial solution satisfies the optimality condition but is not feasible because it violates non-negativity. The algorithm works by iteratively moving through infeasible solutions toward the feasible region while ensuring that each step maintains the optimal objective value relative to the current basis. This approach is particularly useful for re-optimizing after adding constraints.

165. (B) P-2, Q-1, R-4, S-3
Different optimization models are defined by their specific goals and constraints. Assignment problems focus on matching tasks to resources. Transportation models aim to minimize shipping costs across a network. Integer programming involves models where variables must be whole numbers. Finally, PERT and CPM are specialized techniques used for the complex task of project scheduling and management. These mappings clarify the primary application of each technique.

166. (B) P(x) is true for all values of x in the domain.
A universally quantified statement is defined as being true if and only if the predicate it contains holds for every single element within the specified domain. If even one element is found for which the property is false, the entire universal statement is considered false. This rigorous requirement makes universal quantification a powerful tool for expressing absolute mathematical laws and general properties of sets.

167. (D) A set with a relation that is reflexive, anti-symmetric, and transitive.
A partially ordered set, or poset, is a set equipped with a binary relation that satisfies three specific properties: reflexivity, anti-symmetry, and transitivity. This structure allows for a hierarchy where some elements can be compared to others, but not every pair of elements is required to have a defined relationship. This distinguishes posets from totally ordered sets where every pair must be comparable.

NODIA

168. (A) 3/8
When a fair coin is tossed three times, there are eight possible equally likely outcomes. The number of ways to get exactly two heads can be found using the combination formula for picking two positions out of three, which results in three favorable outcomes. Therefore, the probability is three divided by eight. This scenario represents a basic application of binomial distribution principles in probability theory.

169. (C) Used in cryptography and coding theory
Group theory is a branch of abstract algebra that studies the algebraic structures known as groups. A key modern application of this field is in cryptography and coding theory. Concepts such as cyclic groups and the difficulty of the

discrete logarithm problem form the basis for many secure communication protocols, ensuring that information can be transmitted privately and without errors across digital networks.

170. (B) (i), (ii), (iii), (iv)
A general tree is a connected graph that contains no cycles. These two core properties imply several others, such as the existence of a unique path between any two vertices and the fact that the number of vertices is always one more than the number of edges. While a rooted tree has a designated root, this is not a requirement for a general tree.

171. (D) Canonical Product of Sums
The representation of a Boolean function as a product of maxterms is known as the Canonical Product of Sums. In this form, the expression is a conjunction of several sum terms, where each term represents a specific combination of inputs that results in a false output. This is the dual form of the Canonical Sum of Products and provides a standardized way to define functions.

172. (B) The feasible region and optimal solution will remain unchanged.
In linear programming, a redundant constraint is one that does not restrict the feasible region any further than the existing constraints already do. Because it does not change the set of possible solutions, adding such a constraint will have no effect on the shape of the feasible region or the location of the optimal solution. The problem remains mathematically equivalent to its original version.

173. (C) Hypothetical Syllogism
Hypothetical Syllogism is a valid rule of inference that allows for the transitive chaining of conditional statements. It asserts that if p implies q and q implies r, then it must logically follow that p implies r. This rule is essential for constructing long chains of reasoning in deductive proofs and for simplifying complex logical arguments by connecting intermediate steps into a direct conclusion.

174. (D) (A ? B) - (A ∩ B)
The symmetric difference of two sets is a set containing elements that are in either of the sets but not in their intersection. This operation effectively identifies the unique elements that belong to exactly one of the two sets. It can be calculated by taking the union of the two sets and subtracting the elements that they share in common, resulting in the difference.

175. (A) It is used to prove statements about the set of natural numbers.
Mathematical induction is a technique specifically designed to prove properties and statements about the set of natural numbers. By demonstrating that a property holds for an initial case and that its truth for one number guarantees its truth for the next, it establishes the property for the entire infinite sequence of integers. This makes it a fundamental tool in discrete mathematics and number theory.

176. (B) Every field is an integral domain.
The relationship between rings, integral domains, and fields is hierarchical. A field is the most restrictive structure, requiring all ring axioms plus commutativity and the existence of multiplicative inverses for all non-zero elements. Because these requirements encompass all the axioms of an integral domain, every field is automatically an integral domain. However, the reverse is not always true, as seen in integers.

177. (A) An instantaneous code
A prefix code is also known as an instantaneous code because it possesses a unique property that allows for immediate decoding. Since no codeword is a prefix of another, a receiver can identify a complete symbol as soon as the last bit of its codeword arrives, without needing to wait for subsequent bits. this efficiency is vital for real-time data transmission and compression algorithms.

178. (A) Both Statement I and Statement II are true.
Simplifying Boolean functions is a primary goal in digital design to create more efficient hardware. Reducing the number of literals and terms directly translates to a logic circuit with fewer gates and connections. This results in implementations that are not only less expensive to manufacture but also faster and more energy-efficient. Both statements accurately describe the practical motivations behind Boolean algebraic simplification.

Click Here to Buy book on Amazon

179. (C) Unbalanced
An assignment problem is described as unbalanced if the number of available resources, represented by rows, does not equal the number of tasks, represented by columns. Since the standard Hungarian method requires a square cost matrix, dummy rows or columns with zero costs must be added to equalize the dimensions. This ensures that the algorithm can process the problem and find an optimal assignment.

180. (D) Transitive Normal Form (TNF)
Standard normal forms in propositional logic, such as CNF, DNF, and their principal versions, provide specific syntactic structures for representing logical formulas. These are used to standardize expressions for computational analysis. However, there is no recognized structure called Transitive Normal Form in standard logic. Transitivity is a property of relations, but it does not define a standard form for organizing propositional formulas.

181. (C) It is an equivalence relation.
The relation based on absolute values is reflexive because the absolute value of any number equals itself. It is symmetric because if the absolute value of a equals b, then the absolute value of b equals a. Finally, it is transitive because if the absolute values of a and b match, and b and c match, then a and c must also match.

182. (B) 2
There are twelve possible birth months in a year. According to the Pigeonhole Principle, if thirteen people are placed into these twelve monthly holes, at least one month must contain more than one person. Since thirteen is exactly one more than twelve, the principle guarantees that at least two people in the room will share the same birth month, fulfilling the minimum requirement for an overlap.

183. (B) Both A and R are true but R is not the correct explanation of A.
A Hamiltonian circuit is correctly defined as a path that visits every vertex in a graph exactly once and returns to the start. While it is true that many connected graphs do not possess such a circuit, this fact does not provide a logical explanation for the definition. Both statements are true independent facts about graph theory, but they lack causal links.

184. (A) Creating a new structure from two or more given structures, typically on their Cartesian product.
The product of algebraic structures is a method of constructing a new, larger structure from two or more existing ones. This is done by taking the Cartesian product of the underlying sets and defining the algebraic operations component-wise. The resulting structure inherits many properties from its components, allowing mathematicians to build and study complex systems from simpler, well-understood building blocks in a systematic way.

185. (C) Its distance from the root.
In the structural hierarchy of a rooted tree, the level of any given node is defined as its distance from the root node. This distance is measured by counting the number of edges along the unique path connecting the root to that node. The root itself is typically considered to be at level zero, its children at level one, and so forth throughout the tree.

186. (A) x' + y'
To simplify this Boolean expression, one can apply the distributive and complement laws. Factoring out x-prime from the first two terms yields x-prime times the sum of y and y-prime, which simplifies to just x-prime. The remaining expression, x-prime plus x times y-prime, then simplifies further using the distributive property to the final result of x-prime plus y-prime, which is more efficient.

187. (D) To find an initial basic feasible solution for constraints of type '≥' or '='.
Artificial variables are mathematical tools introduced in the Simplex method to provide an initial basic feasible solution for problems with greater than or equal to or equal to constraints. These constraints do not naturally provide a starting point at the origin. The goal of the algorithm is to drive these variables to zero, ensuring that the final solution satisfies all the original problem constraints.

188. (B) $p \rightarrow (p \vee q)$

NODIA

A tautology is a proposition that is always true. In the case of an implication, the statement is true if the antecedent being true guarantees the truth of the consequent. For any proposition p, if p is true, then the disjunction of p and any other variable q must also be true. Therefore, the implication from p to the disjunction of p and q is always true.

189. (C) $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
In propositional logic, operators are evaluated in a specific order to ensure expressions are interpreted consistently. Negation has the highest precedence and is applied first. It is followed by conjunction and then disjunction. Conditional implications come next, and biconditional statements have the lowest priority. This hierarchy, similar to the order of operations in arithmetic, is essential for correctly determining the truth value of complex propositions.

190. (D) 99
A fundamental property of tree graphs is the mathematical relationship between their vertices and edges. Any connected graph with n vertices that contains no cycles must have exactly n minus one edges. Therefore, in a tree consisting of one hundred vertices, there must be exactly ninety-nine edges. This property is used to define and identify trees within the broader field of graph theory.

191. (A) 4, 2, 1, 3
A proof by mathematical induction follows a strict chronological sequence. First, the basis step proves the statement for the initial value. Next, the inductive hypothesis assumes the statement is true for an arbitrary case. Then, the inductive step uses this assumption to prove the property for the next case. Finally, based on these two components, the conclusion generalizes the result for all natural numbers.

192. (B) Associativity
The defining characteristic of a semi-group is that its binary operation must satisfy the property of associativity. This means that when combining three elements, the grouping does not affect the final result. While semi-groups must also satisfy the closure property, associativity is the key structural requirement that distinguishes them from simpler algebraic objects. Other properties like identity or inverse are not required for semi-groups.

NODIA

193. (B) Connectivity and flow
In network analysis and graph theory, the concept of a cut-set is deeply intertwined with connectivity and flow. A cut-set identifies the critical edges that, if removed, would divide the graph and interrupt the flow between different parts of the network. This relationship is formalized in the max-flow min-cut theorem, which equates the maximum possible flow in a network to the minimum capacity cut.

194. (A) To simplify a Boolean function.
The primary application of a Karnaugh Map is the systematic simplification of Boolean functions. By visually representing the truth table in a grid, designers can easily identify groups of adjacent ones that correspond to redundant variables. Eliminating these variables leads to a more compact logical expression, which can be implemented using a smaller and more efficient set of logic gates in digital circuits.

195. (B) (i), (ii), (iv)
PERT and CPM are both critical path management tools with distinct characteristics. PERT is probabilistic, making it suitable for uncertain research projects, while CPM is deterministic and better for predictable construction. Slack measures the flexibility in task timing. However, the critical path is actually the longest path through the network, as it determines the minimum time required to complete the entire sequence of tasks.

196. (D) There is an x such that for all y, P(x,y) is true.
The statement containing an existential quantifier followed by a universal quantifier asserts that there is at least one specific element x that bears the property P to every possible element y in the domain. This is a very strong claim, as it requires a single x to work for all y simultaneously. It logically implies that for any chosen y, one can certainly find an appropriate x.

197. (C) if $(a, b) \in R$ and $(b, c) \in R$, then $(a, c) \in R$
A relation on a set is classified as transitive if it allows the relationship between elements to be transferred through intermediate connections. Specifically, if an element a is related to b, and

that same element b is related to c, then for the relation to be transitive, a must also be directly related to c. This property is essential for defining various types of ordered structures.

198. (B) n^r
When selecting r items from a set of n distinct items where the order of selection matters and items can be reused, the total number of possibilities is found by multiplying n by itself r times. This is because there are n choices available for each of the r positions in the sequence. Mathematically, this is expressed as n raised to the power of r.

199. (D) A ring has two binary operations, while a group has one.
The primary structural difference between a group and a ring is the number of defined binary operations. A group is characterized by a single binary operation that satisfies specific axioms like identity and inverse. In contrast, a ring is a more complex algebraic structure that is equipped with two binary operations, typically addition and multiplication, which interact with each other through the distributive law.

200. (C) P-3, Q-1, R-4, S-2
Different graph concepts are defined by their unique traversal and structural properties. An Eulerian circuit must visit every edge exactly once. A Hamiltonian circuit must visit every vertex exactly once. A spanning tree is a connected, acyclic subgraph that includes all vertices of the parent graph. Finally, a bipartite graph is one that can be colored with exactly two colors.

201. (C) Truth Table
A truth table is a comprehensive and standard way to represent a Boolean function. It lists every possible combination of input variables and the corresponding output for each case. This exhaustive representation clearly defines the behavior of the function, allowing for easy analysis, comparison with other functions, and provides the necessary data for simplification using tools like Karnaugh maps or algebraic laws.

202. (A) Always a convex polygon
In linear programming, the feasible region is formed by the intersection of several linear inequalities. Geometrically, this intersection of half-planes always results in a convex set. For two-dimensional problems, this region takes the shape of a convex polygon. A key property of this shape is that the line segment connecting any two points within the region also lies entirely within the feasible region.

203. (D) Structurally identical, just with different names for elements and operations.
In abstract algebra, an isomorphism is a mapping between two structures that demonstrates they are structurally identical. While the elements and operations in the two structures may have different names or appearances, their underlying behaviors and relationships are exactly the same. An isomorphism preserves the algebraic operations, meaning that any theorem proven for one structure is automatically true for the other isomorphic structure as well.

204. (B) Automated theorem proving using resolution.
Converting logical formulas into Conjunctive Normal Form is a vital preprocessing step for automated theorem proving. The resolution rule, which is the core of many automated reasoning systems, requires the input to be in this specific clause-based format. By standardizing the logic in this way, computers can systematically search for contradictions and prove the validity of complex arguments more efficiently than using traditional methods.

205. (C) The empty relation on a non-empty set is symmetric and transitive.
The empty relation on a non-empty set is considered both symmetric and transitive through vacuous truth. For symmetry, since there are no pairs (a, b) in the relation, the condition that (b, a) must also be present is never violated. Similarly, since there are no sequences of pairs to trigger the transitive condition, it is also never violated. However, it is not reflexive on non-empty sets.

206. (D) The event that is known to have already occurred.
In the mathematical notation for conditional probability, the vertical bar separates the event of interest from the event that is known to have already occurred. In the expression P(A|B), the event B represents the given information or the specific condition that has been satisfied. This known event effectively restricts the sample

Downloaded from www.pdf.tube

space, changing the likelihood of event A based on this new knowledge.

207. (A) Simple path
A path in graph theory is a sequence of edges connecting a sequence of vertices. If the path is restricted such that it never visits the same vertex more than once, it is specifically called a simple path. This distinction is important for many algorithms, such as those finding the shortest path or exploring connectivity, where repeating a vertex would create an unnecessary loop.

208. (B) Standard form
In the context of Boolean algebra and logic, a canonical form is a unique way of representing a function that is standardized. This is also commonly referred to as the standard form. Canonical forms, such as the Sum of Products or Product of Sums, ensure that any two equivalent functions will have the exact same representation, making them essential for comparing and verifying logical expressions.

209. (B) The trade-off between project duration and project cost (crashing).
Cost consideration in project management techniques like PERT and CPM often focuses on the practice of crashing. This involves identifying which activities on the critical path can be shortened by allocating extra resources, thereby reducing the overall project duration. Managers must weigh the increased direct costs of these additional resources against the benefits of finishing the project earlier, such as avoiding penalties or meeting market demands.

210. (D) All (i), (ii), (iii), (iv), and (v) are correct.
The provided statements correctly summarize the hierarchical nature of algebraic structures. A group is indeed a monoid with inverses. A ring combines additive groups and multiplicative semi-groups. Integral domains and fields add further restrictions, such as the absence of zero divisors and the existence of multiplicative inverses respectively. Finally, every field satisfies the ring axioms, though the reverse is not universally true for all rings.

211. (C) Simplification
The rule of inference known as simplification allows one to derive a single part of a conjunction if the entire conjunction is known to be true. If the statement p and q is accepted, it logically follows that p must be true on its own. This rule is frequently used in formal proofs to isolate specific propositions from more complex compound statements for further analysis.

212. (D) A distinguished vertex called the root
A rooted tree is a specialized version of a general tree where one specific vertex is designated as the root. This choice of a starting point imposes a hierarchical structure and a direction on all the edges, leading away from the root toward the leaves. This distinguishes it from an unrooted tree, which is simply a connected graph with no cycles and no orientation.

213. (A) Both A and R are true and R is the correct explanation of A.
The graphical method for solving linear programming problems is restricted to models with a small number of variables because it relies on human visualization. While two-dimensional and some three-dimensional regions can be plotted and analyzed on paper or a screen, humans cannot easily visualize the four-dimensional or higher shapes required for more complex problems. Therefore, the reason accurately justifies the limitation described in the assertion.

NODIA

214. (B) Reflexive, symmetric, and transitive
The relationship is parallel to on a set of lines satisfies all the requirements of an equivalence relation. Every line is parallel to itself, ensuring reflexivity. If line A is parallel to line B, then B is parallel to A, satisfying symmetry. Finally, if A is parallel to B and B to C, then A must be parallel to C, fulfilling the transitivity property.

215. (D) $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|$
To find the total number of unique elements in the union of three sets, the principle of inclusion-exclusion follows a systematic addition and subtraction process. First, the sizes of the individual sets are added. Next, the sizes of all three pairwise intersections are subtracted to correct for double-counting. Finally, the size of the triple intersection is added back to ensure those elements are included exactly once.

216. (A) A graph where multiple edges (parallel edges) are allowed between the same pair of vertices.
A multigraph is a type of graph that relaxes the restriction found in simple graphs regarding edges between vertices. In a multigraph, it is permissible to have multiple edges, also known as parallel edges, connecting the same pair of vertices. This structure is useful for modeling real-world situations where there may be several distinct connections or routes between two points, such as in transportation networks.

217. (C) A + AB = A
The absorption law in Boolean algebra is a simplification rule that shows how a variable can absorb a more complex term. For example, the expression A plus the product of A and B simplifies down to just A. This is because whenever A is true, the entire expression is true regardless of B. This identity is highly effective for reducing the complexity of logical circuits.

218. (A) An infeasible solution
In the Simplex method, artificial variables are used to find an initial solution but should ideally be driven to zero during the optimization process. If the algorithm finishes and one of these variables still has a positive value in the final basis, it indicates that no solution exists that satisfies all the original constraints simultaneously. Thus, the problem is determined to be infeasible.

219. (C) Statement I is true, but Statement II is false.
Every cyclic group is guaranteed to be abelian because any two elements can be expressed as powers of the same generator, and powers of a single element always commute. However, not all abelian groups are cyclic. For example, the Klein four-group is abelian because all its elements commute, but it cannot be generated by a single element, making the second statement false.

220. (B) Root, Left, Right
Pre-order traversal is a depth-first search method for visiting nodes in a tree. It follows a specific pattern where the current node, or root of the subtree, is visited first. After processing the root, the algorithm recursively moves to the left subtree and then to the right subtree. This order is useful for creating a copy of the tree or for prefix notation.

221. (C) Simplifying Boolean expressions
Predicate logic is a powerful tool used for formalizing mathematical statements, expressing natural language arguments, and specifying program properties using quantifiers and variables. However, simplifying Boolean expressions is primarily handled by Boolean algebra and techniques like Karnaugh maps. While predicate logic can represent logical relations, it is not the standard tool used for the specific task of minimizing bitwise Boolean functions in circuit design.

222. (D) Equivalence classes
When an equivalence relation is applied to a set, it groups related elements together into non-overlapping subsets. These disjoint subsets are formally called equivalence classes. Every element of the original set belongs to exactly one of these classes, and any two elements within the same class are related to each other. This structure effectively partitions the set based on the specified equivalence criteria.

223. (A) m - n + 1
To transform a connected graph into a spanning tree, any edges that contribute to cycles must be removed while ensuring the graph remains connected. Since a tree with n vertices must have exactly n minus one edges, the number of edges to be removed from a graph with m edges is simply the difference between m and (n minus one). This equals m minus n plus one.

224. (B) It is event-oriented.
The Program Evaluation and Review Technique is characterized as an event-oriented project management tool. In a PERT network diagram, the focus is placed on the milestones or events that signify the start or completion of activities. This distinguishes it from other methods that may focus more on the activities themselves. PERT also uses probabilistic time estimates to handle uncertainty in complex project schedules.

225. (B) f(a * b) = f(a) • f(b)
A group homomorphism is a function between two groups that preserves the respective group operations. If an operation is performed on two elements in the first group and the result is mapped to the second group, it must yield the same result as mapping the elements

Click Here to Buy book on Amazon

individually and then performing the second group's operation. This structural preservation is the defining characteristic of a homomorphism.

226. (A) P-4, Q-2, R-1, S-3
Accurate terminology is essential in graph theory. The degree of a vertex is the count of edges connected to it. A path is a sequence of connected vertices, while a cycle is a path that returns to its starting point. A simple graph is specifically defined by having no self-loops or parallel edges. Matching these terms to their definitions ensures a consistent mathematical framework.

227. (C) Disjunctive Normal Form
Disjunctive Normal Form is a standard way of writing Boolean expressions as a sum of one or more product terms. This format is also known as Sum of Products. In this form, each term is a conjunction of literals, and these terms are then combined using the disjunction operator. This representation is fundamental in logic design and for specifying the conditions under which functions.

228. (D) It is efficient for sensitivity analysis when a new constraint is added.
The dual simplex method is particularly advantageous when performing sensitivity analysis in linear programming. If a new constraint is added to an already solved problem, the current optimal solution may become infeasible. Instead of starting the entire Simplex process from the beginning, the dual simplex method can efficiently move from this infeasible but optimal point to a new feasible and optimal solution.

229. (B) (i), (ii), (iii), (iv) only
Linear programming models are defined by linear objective functions and constraints, with decision variables typically restricted to non-negative values. A fundamental theorem of linear programming states that if an optimal solution exists, at least one such solution must occur at a corner point, or vertex, of the feasible region. While these models are powerful, they are not directly used for project scheduling without modification.

230. (B) Either true or false, but not both.
In formal logic, a proposition is defined as a declarative sentence that possesses a definite truth value. This means the statement must be either true or false, but it cannot be both at the same time and it cannot be neither. This binary nature allows propositions to be used in logical operations and for constructing the complex arguments that form the basis of mathematical proofs.

231. (D) An equivalence relation
A relation on a set that concurrently satisfies the properties of reflexivity, symmetry, and transitivity is formally known as an equivalence relation. These three properties together ensure that the relation behaves similarly to the standard equality relation. Such relations are fundamental in mathematics because they allow for the partitioning of a set into disjoint classes of elements that are considered equivalent under the specified criteria.

NODIA

232. (A) To show that if the statement holds for an arbitrary case k, it must also hold for the next case k+1.
The inductive step is the dynamic part of a proof by mathematical induction. Its purpose is to establish a logical link between successive cases in a sequence. By demonstrating that the truth of the statement for an arbitrary case k necessarily implies its truth for the next case k plus one, the step ensures that the property propagates infinitely through the entire set of integers.

233. (C) Group
The set of all invertible n by n matrices forms a group under the operation of matrix multiplication, known as the general linear group. This structure satisfies all group axioms: the product of two invertible matrices is invertible, multiplication is associative, the identity matrix serves as the identity element, and every invertible matrix has a unique multiplicative inverse within the same set by definition.

234. (A) Last
Post-order traversal is a depth-first traversal method used for trees and graphs. In this specific algorithm, the traversal processes all the children of a node before finally visiting the node itself. For a binary tree, this means the left subtree is traversed first, then the right subtree, and finally

the root node. Consequently, the root of the entire tree is always visited last.

235. (D) Obtain a logic circuit with the minimum number of gates.
The primary motivation for simplifying Boolean functions in digital logic design is to minimize the hardware required for implementation. A simplified expression contains fewer literals and terms, which directly translates to a circuit with a smaller number of logic gates. This reduction in components leads to lower manufacturing costs, decreased physical size, lower power consumption, and often faster signal propagation through the system.

236. (B) All supply and demand values are equal to 1.
The assignment problem is a specific, restricted version of the more general transportation model. It represents a scenario where each source has a supply of exactly one and each destination has a demand of exactly one. This ensures a one-to-one matching between resources and tasks. Because of this specialized structure, the assignment problem can be solved using efficient algorithms like the Hungarian method.

237. (C) $\neg q \rightarrow \neg p$
The contrapositive of a conditional statement is created by both negating and reversing the antecedent and the consequent. If the original statement is if p then q, its contrapositive is if not q then not p. In logic, a statement and its contrapositive are logically equivalent, meaning they always share the same truth value, which is useful for proving theorems that are difficult directly.

238. (D) Lattice
A lattice is a specialized type of partially ordered set with an additional structural requirement. For any two elements in a lattice, there must exist a unique least upper bound, also known as a join, and a unique greatest lower bound, known as a meet. This property ensures that the set has a well-defined and consistent structure that is more restricted than a general poset.

239. (A) Both A and R are true and R is the correct explanation of A.
A tree is a fundamental graph structure defined by being both connected and acyclic. A well-known theorem in graph theory states that any connected graph with n vertices is a tree if and only if it has exactly n minus one edges. Since the reason correctly identifies the core definition of a tree which leads to the edge count property, both statements are true and linked.

240. (A) Both Statement I and Statement II are true.
Fundamental theorems in group theory prove the uniqueness of certain key elements. For any group, there can be only one identity element that satisfies the identity property for all group members. Furthermore, for each individual element within the group, there exists exactly one unique inverse element that, when combined with it, results in the identity. Both statements accurately reflect these essential algebraic properties.

241. (B) It has zero slack.
In a project network, the critical path consists of the sequence of activities that have no scheduling flexibility. These activities are characterized by having zero slack, meaning any delay in their start or completion will immediately postpone the entire project's finish date. Monitoring these zero-slack tasks is vital for project managers to ensure that the project stays on schedule and meets its required deadline.

242. (A) 16
According to Cayley's formula, the number of spanning trees in a complete graph with n vertices is equal to n raised to the power of n minus two. For the complete graph with four vertices, this calculation is four raised to the power of two, which equals sixteen. This formula provides a quick way to determine the number of possible spanning trees in complete networks.

243. (C) $\neg p \wedge \neg q$
According to De Morgan's Laws, the negation of a disjunction is logically equivalent to the conjunction of the negations of the individual propositions. In this case, saying it is not the case that p or q is true is the same as asserting that p is not true and q is not true. This equivalence is fundamental for simplifying complex logical expressions in mathematics.

244. (D) Prior probability

Downloaded from www.pdf.tube

In the context of Bayes' Theorem, the term P(A) represents the prior probability. This is the initial probability of the hypothesis or event being true before any new evidence or data is taken into account. It is based on previous knowledge or existing beliefs. Bayes' Theorem then uses this prior probability along with new information to calculate the updated, or posterior, probability.

245. (C) with the minimum possible total edge weight.
A minimal spanning tree of a weighted, connected graph is a spanning tree where the total sum of the weights of all its edges is the smallest possible value among all possible spanning trees. While a spanning tree must connect all vertices and contain no cycles, the minimal version focuses on achieving this connectivity with the least total weight, representing the most efficient network.

246. (B) 3, 2, 1
The evolution of project management techniques followed a specific historical progression. The process began with the use of Gantt Charts in the early twentieth century for basic task visualization. Later, in the nineteen-fifties, the more sophisticated Critical Path Method and the Program Evaluation and Review Technique were developed nearly simultaneously to handle the complexities of modern engineering and research projects, in that order.

247. (D) has no zero divisors
An integral domain is a specialized algebraic structure that extends the properties of a commutative ring with unity. The critical additional requirement for a ring to be an integral domain is that it must not contain any zero divisors. This means that if the product of two elements in the set is zero, then at least one of those individual elements must itself be zero.

248. (A) An integer point on or within the feasible region of the corresponding LP problem.
An integer programming problem restricts the decision variables to integer values. Therefore, any valid solution must be an integer point that lies within or on the boundary of the feasible region defined by the problem's constraints. While this region is the same as the one for the corresponding linear programming relaxation, the optimal integer solution is often different from the optimal fractional solution.

249. (C) Every planar graph can be colored with 6 colors.
A famous theorem in graph theory, the Four Color Theorem, states that any planar graph can be colored with at most four colors such that no two adjacent vertices share the same color. While it is true that six colors are more than enough, the statement is technically incorrect because the standard requirement and the specific mathematical bound for planar graphs is actually four.

250. (A) P-4, Q-1, R-3, S-2
Different branches of logic and mathematics focus on distinct concepts. Propositional logic deals with simple statements and truth values using connectives. Predicate logic introduces variables and quantifiers to describe properties. Set theory provides a foundation based on collections of objects. Finally, group theory explores the algebraic properties of structures with a single operation. Matching these fields to their core definitions clarifies their roles.

251. (B) For standardization and simplification of the expressions.
Logical expressions in propositional logic are often transformed into specific structures like Conjunctive Normal Form or Disjunctive Normal Form. These normal forms serve as essential tools for the standardization of complex statements. By utilizing these formats, the underlying logic is streamlined, facilitating a more straightforward simplification of expressions while maintaining their core logical values.

NODIA

252. (A) Propositional logic deals with simple true/false statements, while predicate logic incorporates variables and quantifiers.
Propositional logic focuses on declarative statements categorized as either true or false, using various connectives for combinations. In contrast, predicate logic provides a more robust framework by introducing variables and specific predicates to describe relationships. It also includes universal and existential quantifiers, which allow for more detailed assertions about entire collections of objects or specific entities.

253. (C) They provide a mechanism for generating valid conclusions from given statements.

Logic rules of inference, such as Modus Ponens and Modus Tollens, establish a formal methodology for deductive reasoning. These rules allow for the systematic derivation of new conclusions from a given set of established premises. By applying these structured methods, one can construct logically sound and valid arguments, ensuring that the resulting inferences are consistent with the original data.

254. (D) A statement with mixed quantifiers, like ∀x ∃y, may have a different meaning than ∃y ∀x.
In the context of predicate logic, the sequence of quantification significantly influences the interpretation of a logical statement. When dealing with mixed quantifiers, changing the position of the universal and existential symbols alters the scope and relationship of the variables involved. This demonstrates that the specific arrangement of these elements determines the unique meaning of complex mathematical expressions.

255. (C) Predicates, variables, and quantifiers.
Predicate logic enhances the basic framework of propositional logic by incorporating more detailed structural elements. It introduces predicates to describe properties or relationships and variables to represent subjects. Additionally, the use of universal and existential quantifiers allows for making broad or specific claims about sets of items. These additions collectively provide a more expressive system for advanced logical reasoning.

256. (B) Reflexive, Symmetric, Transitive.
Equivalence relations are fundamental in discrete mathematics for categorizing elements within a set based on specific criteria. To qualify, a relation must simultaneously satisfy three properties: it must be reflexive, meaning every element relates to itself; symmetric, ensuring bidirectional relationships; and transitive. These combined characteristics allow the relation to effectively organize the set into distinct, non-overlapping groups.

257. (D) A partial order uses the antisymmetric property, whereas an equivalence relation uses the symmetric property.
The distinction between equivalence relations and partially ordering relations lies in their specific structural requirements. While both must be reflexive and transitive, they differ in their third property. Equivalence relations require symmetry, facilitating the partitioning of sets. Conversely, partial orders rely on the antisymmetric property, which establishes a hierarchical structure where bidirectional relations only occur when the elements are identical.

258. (A) It divides the set into non-overlapping equivalence classes.
When an equivalence relation is applied to a set, it creates a structural organization known as partitioning. This process groups elements that are related to each other into specific subsets called equivalence classes. These classes are disjoint, meaning no element can belong to more than one group. This systematic division is a direct consequence of satisfying reflexivity, symmetry, and transitivity.

259. (C) The symmetric property, because if $a \leq b$, it does not necessarily mean $b \leq a$.
A relation defined by a less-than-or-equal-to condition fails to be an equivalence relation because it lacks the symmetric property. Although it is reflexive and transitive, the relationship is one-directional for distinct values. If one integer is less than another, the reverse cannot be true. Therefore, the absence of symmetry prevents this hierarchical structure from partitioning the set into classes.

260. (A) There can exist elements 'a' and 'b' in the set for which neither 'a is related to b' nor 'b is related to a' is true.
In a partially ordered set, the concept of comparability dictates whether two specific elements are related under the defined rules. The term implies that for certain pairs, neither element is related to the other in either direction. This characteristic distinguishes partial orders from total orders, where every pair must be comparable. It allows for branches in the hierarchical structure of the set.

261. (D) Combinations.
Counting scenarios often require deciding whether the sequence of selected items is important. In this instance, a committee is formed where the specific order of the professors does

Click Here to Buy book on Amazon

not change the group's composition. This makes combinations the appropriate mathematical tool for the calculation. Combinations focus on the selection of a subset from a larger set without considering the arrangement.

262. (B) To revise the probability of an event given new, related information.
Bayes' Theorem serves as a critical component in conditional probability by facilitating the revision of existing beliefs. It allows for the update of a specific event's probability when new, relevant evidence or data becomes available. By incorporating prior knowledge and conditions related to the event, the theorem provides a mathematically rigorous way to refine likelihood estimates in uncertain environments.

263. (C) Proving a base case and an inductive step.
Mathematical induction is a powerful proof technique used to demonstrate that a statement holds true for all natural numbers. The process consists of two essential phases. First, a base case is established to prove the statement for the initial value. Second, an inductive step shows that if the statement is true for an arbitrary integer, it must follow for the next.

264. (D) The Inclusion-Exclusion Principle.
When calculating the total number of elements in the union of multiple overlapping groups, a specific counting rule is necessary. This principle involves adding the sizes of the individual sets and then subtracting the size of their intersections to prevent double-counting. This systematic approach ensures an accurate count of unique individuals belonging to at least one of the surveyed subject categories.

265. (A) At least one email account must have received more than one email.
The Pigeonhole Principle provides a logical conclusion when distributing a certain number of items into fewer containers. In this scenario, having more emails than available accounts necessitates that at least one account contains multiple messages. This principle does not specify exactly which account or how many items it holds, but it guarantees that the distribution cannot be one-to-one for all items.

266. (B) It must have an identity element and an inverse for each element.
Transforming a semi-group into a group requires the addition of specific algebraic properties. While a semi-group already features an associative binary operation on a set, a group must also include an identity element that leaves others unchanged. Furthermore, every individual element within the set must possess a corresponding inverse element that, when operated upon, returns the structure's defined identity element.

267. (B) An isomorphism is a bijective mapping, while a homomorphism is not necessarily bijective.
Mappings between algebraic structures are classified based on their mathematical properties. A homomorphism is a function that simply preserves the operation between two structures. However, an isomorphism represents a more restrictive relationship. To be considered an isomorphism, the mapping must be a bijection, meaning it is both one-to-one and onto. This indicates that the two algebraic structures are essentially identical in form.

268. (A) It is a mapping from a group onto itself.
Within abstract algebra, an automorphism represents a specific and unique type of isomorphism. While a standard isomorphism maps between two potentially different algebraic structures to show they are identical, an automorphism is a bijective mapping from a group directly onto itself. This process preserves the internal structure and operations of the group while rearranging its elements through a self-contained mapping.

269. (C) Every non-zero element in a field has a multiplicative inverse, a condition not required for an integral domain.
Fields and integral domains are both commutative rings, but they differ in their arithmetic capabilities. An integral domain is characterized by having no zero-divisors. A field goes further by requiring that every non-zero element must have a multiplicative inverse. This additional property makes the field a more versatile structure, allowing for division and more complex arithmetic operations within the algebraic system.

Downloaded from www.pdf.tube

270. (D) H must form a group under the same operation as G.
For a subset of a group to be officially recognized as a subgroup, it must meet specific criteria within the original structure. It is not enough for it to be a simple subset; it must independently satisfy all group axioms. This means the subset must form a complete group while utilizing the exact same binary operation defined by the larger parent group.

271. (C) A weighted graph has numerical values associated with its edges.
Graphs are distinguished by the types of information they convey regarding the connections between vertices. While simple graphs focus solely on the existence of edges, weighted graphs associate specific numerical values with each edge. These weights typically represent costs, distances, or capacities. This addition allows for the analysis of optimization problems, such as determining the most efficient path between two points.

272. (B) An Eulerian circuit visits every edge once, while a Hamiltonian circuit visits every vertex once.
The difference between Eulerian and Hamiltonian circuits lies in the specific components of the graph they are required to visit. An Eulerian circuit must traverse every single edge in the network exactly one time. Conversely, a Hamiltonian circuit focuses on the vertices, requiring a path that visits every vertex exactly once before returning to the start, regardless of the edges used.

273. (D) A salesperson wanting to visit every city on their route exactly once before returning home.
A Hamiltonian path problem involves finding a route that visits every vertex in a given graph exactly once. In a real-world context, this is illustrated by a salesperson who needs to stop at every city on a specific itinerary. Unlike a circuit, the path does not necessarily need to return to the starting point, but the single-visit constraint is essential.

274. (A) A circuit is a path that begins and ends at the same vertex.
In graph theory, the terminology used for traversing edges and vertices depends on the starting and ending locations. A path is defined as a sequence of edges connecting various vertices in a linear fashion. A circuit is a specialized version of a path where the sequence eventually returns to its origin, meaning it starts and ends at the exact same vertex.

275. (C) The sum of the weights of the edges along the path.
Shortest path problems in the context of weighted graphs involve finding the most efficient route between two specific vertices. The optimization goal in these scenarios is to minimize the total cost or distance. This is achieved by calculating and selecting the sequence of edges where the sum of their individual numerical weights results in the lowest possible overall value for the path.

276. (A) It can be drawn on a flat surface without any edges intersecting.
Planarity is a structural property of graphs that determines how they can be represented visually. A graph is defined as planar if it can be embedded or drawn on a two-dimensional surface such that its edges do not cross or intersect one another. This characteristic is particularly important in practical applications like the design of physical networks or printed circuit boards.

277. (D) The graph is divided into two sets, and edges only connect vertices from one set to the other.
Bipartite graphs have a unique structure where the set of vertices is split into two disjoint groups. Within this framework, no edges are permitted to connect vertices that belong to the same set. Instead, every edge must connect a vertex from one group to a vertex in the other. This property makes the entire graph colorable using only two distinct colors.

278. (B) T is a subgraph of G that is a tree and connects all of G's vertices.
A spanning tree is a specific type of subgraph derived from a connected graph. For a subgraph to be considered a spanning tree, it must be a tree, meaning it contains no cycles and remains connected. Crucially, it must also include every single vertex from the original graph. This structure provides a minimal set of edges required to connect all vertices.

NODIA

279. (C) To assign colors to vertices such that no adjacent vertices have the same color.
Graph coloring is a fundamental problem focused on assigning labels to the vertices of a network. The primary objective is to ensure that no two vertices connected by an edge share the same color. This process helps in identifying independent sets of vertices and is used in various scheduling and resource allocation problems where adjacent entities must have distinct assigned values.

280. (D) No codeword can be the beginning part of another codeword.
Prefix codes are specialized sets of codewords used in data compression and information theory. The defining characteristic of such a code is its uniqueness in decoding. This is achieved by ensuring that no codeword in the set serves as the starting sequence, or prefix, for any other codeword. This property allows for instantaneous decoding without the need for additional separators.

281. (A) To reduce the cost and complexity of the corresponding logic circuit.
The process of simplifying Boolean functions is a vital step in the field of digital logic design. By reducing the number of literals and terms in an expression, engineers can create a logically equivalent function that is much easier to implement. This reduction directly leads to the development of logic circuits that are less complex, more efficient, and more economical.

282. (B) The function's output for every possible combination of its inputs.
Truth tables are a foundational method for representing Boolean functions in a clear and exhaustive manner. They provide a complete list of all possible input combinations for the variables involved. For each combination, the table specifies the resulting output value of the function. This comprehensive overview is essential for understanding the function's behavior and for verifying its logical consistency.

283. (B) An expression formed by adding several product terms together.
The Sum of Products form is a standard way to represent Boolean functions algebraically. In this format, the function is constructed by taking several individual product terms and combining them using the logical OR operation. Each product term consists of variables or their complements joined by the AND operation. This structure is commonly used as a starting point for circuit optimization.

284. (A) Binary variables and logical operations.
Boolean algebra provides the mathematical framework necessary for modern computer architecture. It operates primarily on binary variables, which can only take the values of zero or one, representing false and true. These variables are manipulated using fundamental logical operations, including conjunction, disjunction, and negation. Together, these elements form the basis for all logical expressions and digital circuit operations.

285. (C) A standardized representation that is unique for any given function.
Canonical forms in Boolean algebra provide a standardized and unique representation for any given logical function. Unlike simplified expressions, which can take multiple forms, the canonical Sum of Products or Product of Sums is specific to the function's truth table. This uniqueness ensures that two different-looking expressions can be identified as representing the same logic by comparing their canonical formats.

286. (D) Decision variables, objective function, and constraints.
Creating a mathematical model for linear programming involves defining three fundamental elements. Decision variables represent the unknown quantities being calculated. The objective function is a linear equation used to determine the best possible outcome, such as maximizing profit. Constraints are the linear inequalities or equalities that establish the limits and requirements within which the decision variables must operate to be valid.

287. (C) Problems with exactly two decision variables.
The graphical solution method is a visual technique used to solve linear programming problems. However, its practical application is limited to models that involve only two decision variables. This restriction exists because the constraints and the feasible region must be plotted on a two-dimensional coordinate system.

Click Here to Buy book on Amazon

For problems with more variables, more complex algebraic methods like the Simplex algorithm are required.

288. (B) It moves iteratively between vertices of the feasible region to improve the objective function.
The Simplex method is a systematic and efficient algorithm used for solving complex linear programming problems. It operates by identifying a starting vertex on the boundary of the feasible region. The procedure then moves iteratively from one corner point to an adjacent one. With each step, it improves the objective function's value until the optimal solution is eventually reached and confirmed.

289. (D) To determine how the optimal solution changes when model parameters are altered.
Sensitivity analysis is a critical post-optimality procedure in linear programming. Its primary function is to evaluate how changes in the model's parameters affect the final optimal solution. By investigating variations in objective function coefficients or constraint limits, analysts can determine the stability of the solution and understand how real-world fluctuations might impact the best possible outcome of the model.

290. (A) At a corner point of the feasible region.
When using the graphical method to solve a linear programming problem, the feasible region is defined by the intersection of all constraints. A fundamental theorem in optimization states that the optimal solution will always be found at one of the corner points, or vertices, of this region. By testing these specific points, one can efficiently identify the maximum or minimum value.

291. (C) IP problems involve decision variables that are restricted to be integers.
Integer Programming is a specialized branch of optimization that builds upon the foundations of linear programming. The critical distinction between the two lies in the requirements for the decision variables. In Integer Programming, some or all of the variables must result in whole numbers. This restriction is necessary for modeling real-world scenarios where fractional values are physically impossible or meaningless.

292. (A) To minimize the cost of distributing goods from sources to destinations.
The Transportation Model is a specialized linear programming structure designed for logistics and supply chain management. Its fundamental goal is to determine the most cost-effective way to move goods from various sources to multiple destinations. The model accounts for the availability at each source and the requirements at each destination, seeking the distribution plan that minimizes the total shipping cost.

293. (D) All supply and demand values in the model are equal to one.
An assignment model is a refined version of the transportation problem, specifically designed for one-to-one resource allocation. The most distinguishing characteristic of this model is its constraint structure, where the supply at every source and the demand at every destination are exactly equal to one. This ensures that each resource is matched with precisely one task to minimize cost or time.

294. (B) Deciding the number of full-time and part-time staff to hire.
Integer Programming is preferred in scenarios where decision variables represent discrete units that cannot be divided. For example, when determining the number of staff members to hire, a fractional result like 2.5 employees is not actionable. In such cases, the model must enforce integer constraints to ensure that the final solution provides a practical and realistic plan for resource allocation.

295. (C) Supply limits at sources and demand requirements at destinations.
A Transportation Model operates within a specific set of constraints to achieve cost optimization. The primary limitations include supply constraints, which ensure that the amount shipped from any source does not exceed its available capacity. Simultaneously, demand constraints must be satisfied to ensure that each destination receives its required quantity of goods. These rules define the boundaries of the feasible solution.

296. (D) The longest path through the network, determining the minimum project duration.
In project management, the critical path represents the most significant sequence of interdependent activities within a network diagram. It is defined as the longest continuous

NODIA

path from the start to the completion of the project. The total duration of the activities on this specific path determines the minimum amount of time required to finish the entire project without any delays.

297. (A) Any delay in this activity will lead to a delay in the overall project.
The activities located on the critical path are essential for maintaining the project's timeline. Because this path represents the longest duration, these specific tasks have zero slack or float. This means that if any single activity on this path experiences even a minor delay, it will cause a corresponding delay in the final completion date of the overall project.

298. (B) To balance the utilization of resources over the project's duration.
Resource levelling is an optimization technique used to manage the allocation of personnel and materials over a project's lifecycle. Its primary objective is to smooth out the demand for resources, preventing instances of extreme overutilization or idle time. This is achieved by strategically adjusting the start and end dates of non-critical activities without extending the overall project duration.

299. (B) The trade-off between increasing costs and decreasing project duration.
Crashing is an analytical process used in project scheduling to shorten the total duration of a project. This technique involves a careful evaluation of the trade-off between time and money. By allocating additional resources to critical activities, the project can be completed faster, but this acceleration comes at an increased cost. The goal is to find the most economical reduction.

300. (A) A network diagram representing activities and dependencies.
PERT and CPM are essential methodologies for managing complex projects with numerous interconnected tasks. Both techniques share a foundational approach by representing the project as a network diagram. This visual model clearly illustrates the various activities and their logical dependencies. By using this structure, project managers can effectively plan, schedule, and monitor the progress of the entire workflow from start to finish.

www.nodia.in

Downloaded from www.pdf.tube

CHAPTER 02

Computer System Architecture

NODIA

1. What is the fundamental building block of a digital computer that performs a basic logical function?
(A) A register
(B) A logic gate
(C) A flip-flop
(D) A multiplexer

2. Which Boolean algebra property states that A + (B + C) = (A + B) + C?
(A) Associative law
(B) Commutative law
(C) Distributive law
(D) Identity law

3. Given the following statements regarding digital circuits:
Statement I : Combinational circuits have memory elements.
Statement II : The output of a sequential circuit depends only on the present inputs.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

4. Match the Logic Gate in List I with the Boolean Expression in List II.
List I: (Logic Gate)
P. NAND
Q. NOR
R. XOR
S. XNOR
List II: (Boolean Expression)
1. $A \oplus B$
2. $(A + B)'$
3. $(A . B)'$
4. $(A \oplus B)'$
Choose the correct option from those given below:
(A) P-3, Q-2, R-1, S-4
(B) P-2, Q-3, R-1, S-4
(C) P-3, Q-2, R-4, S-1
(D) P-1, Q-2, R-3, S-4

5. Which of the following statements about Karnaugh Map (K-Map) simplification is correct?
(A) Grouping can only be done for pairs of 1s, not for quads or octets.
(B) Don't care conditions cannot be used for grouping to simplify an expression.
(C) A larger group of 1s leads to a simpler term in the final expression.
(D) Each group must contain at least one 1 that is not part of any other group.

6. If a decoder has 'n' input lines, how many possible output lines can it have at maximum?
(A) n
(B) 2^n
(C) 2n
(D) n^2

7. All of the following are types of flip-flops EXCEPT:
(A) SR Flip-Flop
(B) D Flip-Flop
(C) T Flip-Flop
(D) AND Flip-Flop

8. Consider the following statements regarding sequential circuits:
(a) Their output depends on both present inputs and past outputs.
(b) They contain memory elements to store state information.
(c) Asynchronous sequential circuits do not use a clock signal.
(d) A register is an example of a combinational circuit.
Which of the above statement(s) is/are correct?
(A) (a) and (d) only
(B) (a), (b) and (c) only
(C) (b) and (d) only
(D) (a), (b), (c) and (d)

9. A multiplexer with 'n' selection lines can select one out of how many input lines?
(A) n
(B) 2n
(C) 2^n
(D) n^2

10. What is a register in the context of digital circuits?
(A) A group of flip-flops used to store a binary word.
(B) A circuit that selects one of many inputs.
(C) A circuit that converts binary information to an alphanumeric character.
(D) A circuit that performs arithmetic operations.

11. Which of the following statements about counters is incorrect?
(A) A synchronous counter has all its flip-flops triggered by the same clock pulse.
(B) A ripple counter is a type of asynchronous counter.
(C) A counter is a register that goes through a predetermined sequence of states.
(D) Asynchronous counters are faster than synchronous counters for large numbers of bits.

12. Arrange the following steps for designing a combinational logic circuit in the correct chronological order:
(a) Draw the logic diagram.
(b) Obtain the truth table from the problem statement.
(c) Obtain the simplified Boolean functions.
(d) Derive the Boolean functions from the truth table.
Choose the correct option:
(A) a, b, d, c
(B) b, d, c, a
(C) c, b, a, d
(D) b, a, c, d

13. Which of the following best describes the function of a memory unit?
(A) To perform arithmetic and logic operations.
(B) To control the timing and sequence of operations.
(C) To store binary information in the form of bits.
(D) To transfer data between the CPU and peripheral devices.

14. What is the primary cause for the increased circuit complexity and reduced size of modern digital computers?
(A) The use of vacuum tubes
(B) The invention of the transistor
(C) The adoption of Boolean algebra
(D) The development of integrated circuits (ICs)

15. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : The Boolean expression A + A'B is equivalent to A + B.
Reason (R) : This is due to the absorption law, which states X + X'Y = X + Y.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

16. A logic circuit is required to produce a HIGH output only when its three inputs A, B, and C are such that a majority of the inputs are HIGH. Which Boolean expression correctly represents this circuit?
(A) A + B + C
(B) AB + BC + AC
(C) ABC
(D) A'B'C'

17. Identify the correct pair of a logic gate and its complementary gate.
(A) AND - NOR
(B) AND - NAND
(C) OR - NAND
(D) XOR - AND

18. Given the following statements regarding digital logic components:
Statement I : A decoder converts a binary code into a set of signals, activating only one output line.
Statement II : A multiplexer selects one of

Click Here to Buy book on Amazon

several input signals and forwards it to a single output line.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

19. Consider the following statements regarding flip-flops:
(a) A flip-flop is a bistable multivibrator.
(b) The JK flip-flop has a race-around condition.
(c) A D flip-flop can be made from an SR flip-flop by connecting S and R with an inverter.
(d) The T flip-flop is known as the toggle flip-flop.
(e) A flip-flop is a combinational circuit.
Which of the above statement(s) is/are correct?
(A) (a), (b), and (e) only
(B) (c), (d), and (e) only
(C) (a), (b), (c), and (d) only
(D) (a), (c), and (d) only

20. What is the purpose of an 'Integrated Circuit' (IC) in digital electronics?
(A) To function as a dedicated power supply source for integrated systems.
(B) To provide a protective physical housing unit for a single transistor.
(C) To serve as an external storage device similar to a traditional hard disk.
(D) To miniaturize and combine a large number of components into a single chip.

21. A digital computer is defined as a machine that can:
(A) Only perform mathematical calculations.
(B) Process continuous analog signals.
(C) Solve problems by processing information in discrete form.
(D) Function without a memory unit.

22. All of the following are examples of combinational circuits EXCEPT:
(A) Adder
(B) Counter
(C) Decoder
(D) Multiplexer

23. Match the Circuit in List I with its Primary Function in List II.
List I: (Circuit)
P. Full Adder
Q. Multiplexer
R. Decoder
S. Register
List II: (Primary Function)
1. Data selection
2. Code conversion
3. Information storage
4. Arithmetic sum of three bits
Choose the correct option from those given below:
(A) P-4, Q-1, R-2, S-3
(B) P-1, Q-4, R-2, S-3
(C) P-4, Q-2, R-1, S-3
(D) P-3, Q-1, R-2, S-4

24. If a 4-bit binary ripple counter starts from the state 0000, what will be its state after 11 clock pulses?
(A) 1011
(B) 1010
(C) 1100
(D) 1101

NODIA

25. Which of the following statements about number systems is incorrect?
(A) The base or radix of the binary number system is 2.
(B) A hexadecimal digit can be represented by 4 binary bits.
(C) The octal number system uses digits from 0 to 8.
(D) The decimal number 10 is represented as 'A' in hexadecimal.

26. In data representation, what does the term 'Fixed Point Representation' refer to?
(A) A representation where the position of the radix point is fixed.
(B) A representation where the decimal point is always at the end of the number.
(C) A representation used exclusively for integers.
(D) A representation that uses a sign, mantissa, and exponent.

27. What is the decimal equivalent of the binary number 10110.01?
(A) 22.5
(B) 21.25

(C) 21.5
(D) 22.25

28. Given the following statements regarding number complements:
Statement I : The 2's complement of a binary number is obtained by taking the 1's complement and adding 1.
Statement II : The primary advantage of 2's complement representation is that subtraction can be performed by addition.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

29. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : Floating-point representation can represent a much wider range of numbers compared to fixed-point representation with the same number of bits.
Reason (R) : Floating-point representation separates the number into a mantissa and an exponent, allowing the radix point to float.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

30. In a fixed-point representation using 'n' bits, what does the leftmost bit typically represent?
(A) The least significant digit
(B) The position of the radix point
(C) An exponent value
(D) The sign of the number

31. Which statement correctly describes the purpose of an error detection code like a parity bit?
(A) It can detect the occurrence of an odd number of bit errors.
(B) It can correct any single-bit error in the transmitted data.
(C) It guarantees the data is always transmitted correctly.
(D) It can correct multi-bit errors.

32. Consider the following statements regarding the process of binary addition:
(a) $0 + 0 = 0$
(b) $0 + 1 = 1$
(c) $1 + 1 = 0$ with a carry of 1
(d) $1 + 1 + 1 = 1$ with a carry of 1
Which of the above statement(s) is/are correct?
(A) (a) and (b) only
(B) (a), (b) and (c) only
(C) (c) and (d) only
(D) (a), (b), (c) and (d)

33. For scientific calculations involving very large or very small numbers, which data representation is most appropriate?
(A) Fixed-point representation
(B) Floating-point representation
(C) BCD representation
(D) 2's complement integer representation

34. To perform the subtraction A - B using 2's complement arithmetic, which of the following operations is performed?
(A) A + (2's complement of B)
(B) (2's complement of A) + B
(C) A + (1's complement of B)
(D) (1's complement of A) + (1's complement of B)

35. Which of the following is an incorrect statement regarding the components of a floating-point number?
(A) The mantissa represents the significant digits of the number.
(B) The exponent indicates the position of the radix point.
(C) The sign bit applies only to the exponent.
(D) The base is usually implicit and is not stored.

36. Match the Error Code in List I with its Primary Characteristic in List II.
List I: (Error Code)
P. Parity Bit
R. Hamming Code
S. CRC (Cyclic Redundancy Check)
T. Checksum
List II: (Primary Characteristic)
1. Corrects single-bit errors and detects double-bit errors.

NODIA

2. Detects single-bit errors.
3. Based on polynomial division.
4. Based on sum of data blocks.

Choose the correct option from those given below:
(A) P-1, R-2, S-3, T-4
(B) P-2, R-1, S-3, T-4
(C) P-2, R-3, S-1, T-4
(D) P-3, R-1, S-2, T-4

37. Which of the following is not a standard data type directly handled by the arithmetic processing unit of a computer?
(A) Fixed-point numbers
(B) Floating-point numbers
(C) Character strings
(D) Binary coded decimal (BCD)

38. In fixed-point addition, an overflow occurs when:
(A) The sum of two positive numbers results in a negative number.
(B) The sum requires more bits than the allocated storage.
(C) The sum of a positive and a negative number is zero.
(D) The result is a non-integer value.

39. Given the following statements regarding number representations:
Statement I : In 1's complement, there are two representations for zero (+0 and -0).
Statement II : In 2's complement, there is a unique representation for zero.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

40. Consider the following statements regarding number systems:
(a) The hexadecimal system uses 16 distinct symbols.
(b) The binary equivalent of octal 75 is 111101.
(c) The decimal number system is a positional number system.
(d) The process of converting from decimal to another base involves division by the radix.
(e) The binary number system is not positional.
Which of the above statement(s) is/are correct?
(A) (a), (b), and (e) only
(B) (a), (b), (c), and (d) only
(C) (c), (d), and (e) only
(D) (a) and (d) only

41. What is the main advantage of Hamming codes over simple parity codes?
(A) They are significantly simpler to implement within hardware logic.
(B) They use fewer redundant bits than a standard single-parity code.
(C) They can detect the position of a single-bit error and correct it.
(D) They are restricted to being used only for strictly binary data sets.

42. If a computer uses an 8-bit 2's complement representation for signed integers, what is the decimal value of the binary number 11110101?
(A) -11
(B) 245
(C) -10
(D) -245

43. What does Register Transfer Language (RTL) primarily describe?
(A) The physical wiring of a computer.
(B) High-level programming language constructs.
(C) The structure of the memory hierarchy.
(D) The sequence of microoperations and the registers involved.

44. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : Error detection codes like CRC are essential for reliable data transmission and storage.
Reason (R) : Noise and physical defects in the communication channel or storage medium can cause data to be altered.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

Downloaded from www.pdf.tube

45. Identify the correct pair representing a computer arithmetic algorithm and its purpose.
(A) Booth's algorithm - Floating-point addition
(B) Restoring division - Signed integer multiplication
(C) Booth's algorithm - Signed integer multiplication
(D) Non-restoring division - Unsigned integer addition

46. In the context of computer arithmetic, Booth's algorithm for multiplication is advantageous because it can:
(A) Multiply two numbers in a single clock cycle.
(B) Work directly on floating-point numbers without conversion.
(C) Eliminate the need for an accumulator register.
(D) Handle multiplication of signed binary numbers efficiently.

47. Which of the following statements regarding number system conversion is incorrect?
(A) To convert a decimal integer to binary, the number is repeatedly multiplied by 2.
(B) To convert a binary number to octal, bits are grouped in threes starting from the LSB.
(C) To convert a binary number to hexadecimal, bits are grouped in fours starting from the LSB.
(D) To convert a decimal fraction to binary, the fraction is repeatedly multiplied by 2.

48. The standard division algorithm in digital computers typically involves a sequence of which operations?
(A) Addition and logical AND
(B) Comparison, shifting, and subtraction
(C) Multiplication and shifting
(D) Negation and addition

49. What is a microoperation?
(A) An instruction written in a high-level language.
(B) An elementary operation performed on data stored in registers.
(C) The process of fetching an instruction from memory.
(D) A large-scale operation involving the entire memory.

50. Given the following statements regarding data coding:
Statement I : Gray code is a weighted code.
Statement II : Binary Coded Decimal (BCD) represents each decimal digit with a 4-bit binary code.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

51. The notation `R2 ← R1` in Register Transfer Language implies that:
(A) The contents of R1 and R2 are swapped.
(B) The contents of R1 are added to the contents of R2.
(C) The contents of R1 are transferred to R2.
(D) The contents of R2 are transferred to R1.

52. What is the primary purpose of a common bus system in a digital computer?
(A) To provide a shared communication path between multiple registers and units.
(B) To increase the operational speed of internal arithmetic operations.
(C) To store the program instructions permanently in non-volatile memory.
(D) To decode instructions and generate the necessary control signals.

53. Given the following statements regarding microoperations:
Statement I : Arithmetic microoperations perform arithmetic on numeric data in registers.
Statement II : Logic microoperations perform bitwise manipulations on non-numeric data in registers.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

54. Match the Microoperation in List I with its Symbolic Representation in Register Transfer Language (RTL) in List II.
List I: (Microoperation)
(P) Arithmetic Shift Right
(Q) Logical Shift Left
(R) Complement

NODIA

(S) Clear
List II: (Symbol)
(1) `shl`
(2) `shr`
(3) `R1 ← 0`
(4) `R1 ← R1''
Choose the correct option from those given below:
(A) P-1, Q-2, R-4, S-3
(B) P-2, Q-1, R-3, S-4
(C) P-2, Q-1, R-4, S-3
(D) P-1, Q-2, R-3, S-4

55. Which statement correctly describes a memory read operation using RTL?
(A) It is denoted by `M[AR] ← DR`.
(B) It involves transferring data from a CPU register to memory.
(C) It is denoted by `DR ← M[AR]`.
(D) It requires only the data register (DR) to be specified.

56. If a register R1 contains the binary value 1100 and R2 contains 1010, what is the content of R1 after the microoperation `R1 ← R1 ⊕ R2` (XOR) is executed?
(A) 0110
(B) 1100
(C) 1010
(D) 1000

57. The three types of shift microoperations are logical, circular, and:
(A) Horizontal
(B) Vertical
(C) Boolean
(D) Arithmetic

58. Consider the following statements about a bus system for register transfers:
(a) A bus system is created using multiplexers or three-state buffers.
(b) It reduces the number of connections required compared to direct wiring.
(c) Only one transfer can take place on the bus at any given time.
(d) The control signals determine which register is selected for the transfer.
Which of the above statement(s) is/are correct?
(A) (a) and (d) only
(B) (b) and (c) only
(C) (a), (b), and (c) only
(D) (a), (b), (c), and (d)

59. What is a microoperation?
(A) The process of assembling a program.
(B) An elementary operation performed during one clock pulse on data in registers.
(C) A complete instruction executed by the CPU.
(D) A data transfer between the CPU and I/O devices.

60. Which of the following is an incorrect representation in Register Transfer Language?
(A) `T: R2 ← R1` (A conditional transfer)
(B) `R1 + R2` (Represents the addition operation)
(C) `M[AR] ← R3` (A memory write operation)
(D) `AR ← AR + 1` (Incrementing a register)

61. Which RTL expression most appropriately represents the statement If P=1, then transfer the contents of R1 to R2?
(A) `R2 ← R1(P)`
(B) `P(R2 ← R1)`
(C) `P: R2 ← R1`
(D) `R1 → R2 | P=1`

62. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : A common bus system is used for transferring information between registers in a multiple-register configuration to reduce the number of wires.
Reason (R) : A bus structure connects the source and destination registers directly for every possible transfer, increasing complexity.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

63. The use of a common bus system in a CPU leads to a reduction in circuit complexity. What is the main trade-off for this simplification?
(A) Increased power consumption.

Click Here to Buy book on Amazon

(B) Slower overall transfer speed due to shared access.
(C) Inability to perform arithmetic operations.
(D) Requirement for a larger number of registers.

64. An arithmetic circuit with 'n' bits can be constructed with:
(A) n multiplexers.
(B) n decoders.
(C) n full-adders.
(D) n flip-flops.

65. Given the following statements regarding shift operations:
Statement I : A logical shift-right operation fills the leftmost bit position with a 0.
Statement II : An arithmetic shift-right operation preserves the sign bit by filling the leftmost position with the original sign bit.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

66. Consider the following microoperations:
(a) AND
(b) Add
(c) OR
(d) Complement
(e) Shift
Which of the above are considered logic microoperations?
(A) (b) and (e) only
(B) (a), (c), and (d) only
(C) (a), (b), (c), and (d) only
(D) (a), (c), (d), and (e) only

67. A 4-bit register R contains the value 1011. What will be the value in R after an arithmetic shift right operation?
(A) 0101
(B) 1101
(C) 1100
(D) 0011

68. Identify the correct pair of a microoperation and its description.
(A) Selective-complement - Complements all bits in register A.
(B) Mask - Clears all bits in register A.
(C) Insert - Sets bits in register A where corresponding bits in B are 0.
(D) Selective-set - Sets bits in register A where corresponding bits in B are 1.

NODIA

69. In the RTL notation `P: R1 ← R2`, what does 'P' represent?
(A) A processor
(B) A program
(C) A control function or signal
(D) A peripheral device

70. What is the primary function of a three-state bus buffer?
(A) To store a binary bit of information.
(B) To act as a switch that can disconnect a device from the bus.
(C) To convert analog signals to digital.
(D) To amplify the signal on the bus.

71. The Stored Program Organization concept, attributed to von Neumann, implies that:
(A) Programs and data are stored in separate memories.
(B) Only programs, not data, are stored in memory.
(C) Instructions and data are stored together in the same memory.
(D) The computer can only execute programs stored on external media.

72. In a basic computer, which register holds the address of the next instruction to be executed?
(A) Accumulator (AC)
(B) Instruction Register (IR)
(C) Program Counter (PC)
(D) Address Register (AR)

73. Given the following statements regarding instruction codes:
Statement I : An instruction code is a group of bits that instruct the computer to perform a specific operation.
Statement II : The opcode is the part of the instruction code that specifies the operands for the operation.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.

(D) Statement I is false, but Statement II is true.

74. Arrange the basic phases of an instruction cycle in the correct chronological order:
(a) Decode the instruction.
(b) Fetch the instruction from memory.
(c) Execute the instruction.
(d) Read the effective address from memory if the instruction is indirect.
Choose the correct option:
(A) b, a, d, c
(B) a, b, c, d
(C) b, d, a, c
(D) d, b, a, c

75. What is the role of the Timing and Control unit in a computer?
(A) To perform arithmetic calculations.
(B) To store data temporarily.
(C) To generate timing signals for sequencing all operations.
(D) To hold the address of the active memory location.

76. If a memory-reference instruction uses indirect addressing, what does the address part of the instruction specify?
(A) The operand itself.
(B) The address of the operand.
(C) The address of the memory word that contains the address of the operand.
(D) The register that contains the operand.

77. All of the following are types of instructions in a basic computer EXCEPT:
(A) Memory-reference instructions
(B) Register-reference instructions
(C) Input-output instructions
(D) Floating-point instructions

78. Consider the following statements regarding the interrupt cycle:
(a) It is initiated by an external or internal signal.
(b) The address of the current program is stored in a temporary location.
(c) The program counter is loaded with the starting address of the interrupt service routine.
(d) It is a hardware-implemented cycle.
Which of the above statement(s) is/are correct?
(A) (a) and (d) only
(B) (a), (b) and (c) only
(C) (b) and (d) only
(D) (a), (b), (c) and (d)

79. What is the most appropriate function of the Input-Output and Interrupt system?
(A) To execute program code and manage software tasks.
(B) To manage the information flow between the computer and external devices.
(C) To perform floating-point arithmetic on complex numbers.
(D) To store the operating system kernel in permanent storage.

80. The instruction cycle is the sequence of steps that the CPU goes through to:
(A) Process an interrupt.
(B) Start up the computer.
(C) Process one instruction.
(D) Access data from the hard disk.

81. Which of the following statements about the instruction format of a basic computer is typically incorrect?
(A) It is divided into an address part and an opcode part.
(B) The opcode part specifies the operation to be performed.
(C) A mode bit is often included to specify direct or indirect addressing.
(D) All instructions must have an address part.

82. Match the Register in List I with its Primary Function in List II.
List I: (Register)
(P) AR (Address Register)
(Q) PC (Program Counter)
(R) DR (Data Register)
(S) IR (Instruction Register)
List II: (Function)
(1) Holds instruction code
(2) Holds memory operand
(3) Holds address of instruction
(4) Holds address for memory
Choose the correct option from those given below:
(A) P-4, Q-3, R-2, S-1
(B) P-3, Q-4, R-2, S-1
(C) P-4, Q-3, R-1, S-2

Downloaded from www.pdf.tube

(D) P-1, Q-2, R-3, S-4

83. When an interrupt occurs, the CPU suspends the execution of the current program. This is primarily caused by:
(A) The program voluntarily yielding control.
(B) A hardware signal from a peripheral device or an internal event.
(C) The operating system deciding to switch tasks.
(D) The program encountering a fatal error.

84. What register in the basic computer holds the instruction that is currently being executed?
(A) Program Counter (PC)
(B) Instruction Register (IR)
(C) Accumulator (AC)
(D) Data Register (DR)

85. Given the following statements regarding addressing modes:
Statement I : In direct addressing, the address field of the instruction contains the operand itself.
Statement II : In indirect addressing, the address field of the instruction contains the address of the operand.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

86. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : The stored program concept is fundamental to modern general-purpose computers.
Reason (R) : Storing instructions in memory allows the computer to manipulate the instructions, enabling features like loops and subroutines easily.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

87. A computer has a memory-reference instruction ADD X, where X is a memory address. If the addressing mode is direct, what does the CPU do to get the operand?
(A) It treats the value 'X' as the operand.
(B) It fetches the data from the memory location specified by X.
(C) It fetches the data from the memory location whose address is stored at location X.
(D) It fetches the data from the register specified by X.

88. Identify the correct pair of a basic computer instruction and its type.
(A) CLA (Clear Accumulator) - Memory-reference
(B) AND (Bitwise AND) - Register-reference
(C) INP (Input Character) - Input-Output
(D) BUN (Branch Unconditionally) - Register-reference

NODIA

89. What is the 'effective address' in the context of computer instructions?
(A) The address of the instruction itself.
(B) The physical address of the operand.
(C) The address of the next instruction to be fetched.
(D) The address specified in the address field of the instruction.

90. Consider the following components of a basic computer:
(a) A memory unit to store instructions and data.
(b) A processing unit to execute instructions.
(c) A control unit to interpret instructions and generate control signals.
(d) A set of registers for temporary storage.
(e) A high-speed network interface.
Which of the above components are essential for the operation of a basic stored-program computer?
(A) (a), (b), and (e) only
(B) (c), (d), and (e) only
(C) (a), (b), (c), and (d) only
(D) (a), (c), and (d) only

91. The fetch cycle in a basic computer involves which of the following microoperations?
(A) $AR \leftarrow PC$, $IR \leftarrow M[AR]$, $PC \leftarrow PC + 1$
(B) $IR \leftarrow M[AR]$, $PC \leftarrow PC + 1$

(C) AR ← PC, PC ← PC + 1
(D) M[AR] ← DR, PC ← PC + 1

92. Which statement about memory-reference instructions in a basic computer is incorrect?
(A) They require an address part to specify a memory operand.
(B) They use an opcode to specify the operation (e.g., ADD, LDA).
(C) They can have either direct or indirect addressing.
(D) They operate only on data present in CPU registers.

93. During an interrupt cycle, what is the primary reason for saving the content of the Program Counter (PC)?
(A) To allow the interrupt service routine to modify it.
(B) To enable the CPU to return to the correct point in the interrupted program.
(C) To free up the PC for use by the I/O device.
(D) To check for program errors.

94. Which of the following is the most appropriate description of a computer instruction?
(A) A binary code specifying a sequence of microoperations for the computer.
(B) A high-level language statement like `if-then-else`.
(C) A hardware signal that triggers an operation.
(D) A piece of data stored in the accumulator.

95. In a basic computer's input-output configuration, what is the role of the FGI (Input Flag) and FGO (Output Flag) flags?
(A) To store the data being transferred.
(B) To indicate an error in the I/O device.
(C) To synchronize the timing of data transfer between the CPU and the I/O device.
(D) To select which I/O device is active.

96. What is Assembly Language?
(A) A low-level symbolic language that has a one-to-one correspondence with machine language.
(B) A high-level language that is portable across different computer architectures.
(C) A language that uses binary codes to represent instructions.
(D) A language used exclusively for writing operating systems.

97. The primary function of an assembler is to:
(A) Execute a program written in a high-level language.
(B) Translate a program from assembly language to machine language.
(C) Link several object files into a single executable file.
(D) Interpret and execute a program one instruction at a time.

98. Given the following statements regarding computer languages:
Statement I : Machine language is directly executable by the computer's processor.
Statement II : Assembly language uses mnemonics for instructions, making it more human-readable than machine language.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

99. Consider the following tasks related to an assembler:
(a) Generate the machine code for instructions.
(b) Create a symbol table with addresses for all user-defined labels.
(c) Process pseudo-instructions like ORG and END.
(d) Resolve forward references.
Which of the above tasks are typically performed during the first pass of a two-pass assembler?
(A) (a) and (d) only
(B) (b) and (c) only
(C) (a), (b) and (c) only
(D) (b), (c) and (d) only

100. Which statement correctly describes a program loop in assembly language?
(A) A section of code that is executed only once during execution.
(B) A hardware component that repeats a specific instruction set.
(C) A sequence of instructions executed repeatedly until a condition is met.
(D) A type of subroutine that cannot be called from the main program.

Click Here to Buy book on Amazon

101. To implement a subroutine call, which instruction is most crucial for ensuring a proper return to the calling program?
(A) An instruction that stores the return address.
(B) An instruction that clears the accumulator.
(C) An instruction that halts the processor.
(D) An instruction that performs an unconditional jump.

102. In assembly language programming, all of the following are examples of executable instructions EXCEPT:
(A) ADD
(B) LDA
(C) ORG
(D) JMP

103. Match the Assembly Language Directive in List I with its Purpose in List II.
List I: (Directive)
(P) ORG
(Q) END
(R) DEC
(S) HEX
List II: (Purpose)
(1) Signifies the end of the program
(2) Declares a decimal number
(3) Declares a hexadecimal number
(4) Sets the origin (starting address) of the program
Choose the correct option from those given below:
(A) P-1, Q-4, R-2, S-3
(B) P-4, Q-1, R-3, S-2
(C) P-4, Q-1, R-2, S-3
(D) P-3, Q-1, R-2, S-4

104. What is the most appropriate reason for using subroutines in programming?
(A) To increase the overall size of the binary program code.
(B) To make the program run faster by eliminating logical jumps.
(C) To avoid code repetition and structure the program into modular blocks.
(D) To allow direct manipulation of all underlying hardware registers.

105. What is a symbol table, as used by an assembler?
(A) A table that stores the machine code for all instructions.
(B) A table that maps symbolic labels to their memory addresses.
(C) A table of hardware registers and their current values.
(D) A table that stores the results of arithmetic operations.

106. Which of the following statements about Input-Output programming is incorrect?
(A) It involves checking status flags to synchronize data transfer.
(B) It can be implemented using interrupts.
(C) It does not require any special instructions; it uses standard data movement instructions.
(D) The CPU must be fully dedicated to waiting in a programmed I/O loop.

107. The translation from assembly language to machine language is necessary because:
(A) Assembly language is too slow for modern computers.
(B) The CPU can only understand and execute instructions in binary machine code format.
(C) Assembly language programs are too large to fit in main memory.
(D) High-level languages cannot be compiled directly to assembly language.

108. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : Programming in assembly language generally results in faster and more memory-efficient code than programming in a high-level language.
Reason (R) : Assembly language provides direct control over the CPU's registers and operations, allowing for fine-tuned optimization.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

109. The second pass of a two-pass assembler is responsible for:
(A) Building the symbol table.
(B) Finding syntax errors in the source code.

Downloaded from www.pdf.tube

(C) Translating instructions into machine code and resolving symbolic addresses using the symbol table.
(D) Assigning the initial memory location for the program.

110. Given the following statements regarding program interrupts:
Statement I : A program interrupt transfers control to a specific service routine, similar to how a subroutine call transfers control.
Statement II : Unlike a subroutine call, a program interrupt is initiated by an external or internal hardware signal rather than by an instruction in the program.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

111. If a program loop starts at address 200 and needs to execute 10 times, and the loop body is 5 memory words long, a counter-based implementation would typically involve:
(A) Setting a counter to 1 and incrementing it until it reaches 10.
(B) Jumping back to address 200 exactly 10 times.
(C) Storing the loop body in 10 different memory locations.
(D) Setting a register to -10 and incrementing it until it becomes zero.

112. Identify the correct pair matching a programming language to its level.
(A) C++ - Low-level language
(B) Machine Code - High-level language
(C) Assembly Language - Low-level language
(D) Java - Machine language

113. Machine language consists of instructions in what format?
(A) Symbolic mnemonics
(B) English-like statements
(C) Binary or hexadecimal code
(D) Flowchart diagrams

114. The linkage between a calling program and a subroutine is often handled by:
(A) The assembler.
(B) The operating system.
(C) A dedicated hardware stack.
(D) The programmer manually managing addresses.

115. Which statement regarding Input-Output (I/O) programming is incorrect?
(A) In programmed I/O, the CPU continuously checks the status of the I/O device.
(B) Interrupt-driven I/O allows the CPU to perform other tasks while waiting for the I/O operation to complete.
(C) Programmed I/O is the most efficient method for high-speed data transfers.
(D) I/O instructions are needed to check flags and transfer data.

NODIA

116. What is stored in the Control Memory of a microprogrammed control unit?
(A) The user program and data.
(B) The operating system kernel.
(C) The microprogram consisting of microinstructions.
(D) The contents of the CPU registers.

117. Which of the following is a key capability required for address sequencing in a microprogrammed control unit?
(A) Incrementing the control address register.
(B) Loading the control address register from memory.
(C) Performing arithmetic operations on data registers.
(D) Accessing peripheral devices directly.

118. Given the following statements regarding computer control units:
Statement I : A hardwired control unit is implemented using fixed logic circuits like gates, flip-flops, and decoders.
Statement II : A microprogrammed control unit is generally slower than a hardwired control unit.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

119. Consider the following fields in a microinstruction:

(a) A condition field for branching.
(b) A branch address field.
(c) Multiple control fields, one for each distinct microoperation.
(d) The entire user program instruction.
(e) An address for the next microinstruction.
Which of the above are typically found as fields in a horizontal microinstruction?
(A) (d) and (e) only
(B) (a), (b), and (c) only
(C) (a), (c), and (e) only
(D) (a), (b), (c) and (e)

120. Which statement correctly describes the design of a microprogrammed control unit?
(A) It involves designing a complex state machine with combinational logic.
(B) It involves writing a microprogram (firmware) that implements the computer's instruction set.
(C) It is inflexible and cannot be easily modified to support new instructions.
(D) It is primarily used in simple, low-cost processors.

121. In a microprogrammed control unit, the process of fetching and executing a microinstruction is called:
(A) The instruction cycle.
(B) The micro-cycle.
(C) The fetch-decode-execute cycle.
(D) The address sequencing cycle.

122. A microinstruction format typically includes all of the following fields EXCEPT:
(A) Control word bits to enable microoperations.
(B) Condition bits for branching.
(C) Branch address field.
(D) Main memory address for data operand.

123. Match the Term in List I with its Description in List II.
List I: (Term)
(P) Control Address Register (CAR)
(Q) Control Memory
(R) Microinstruction
(S) Microprogram Sequencer
List II: (Description)
(1) A control word that specifies one or more microoperations.
(2) A unit that generates the address of the next microinstruction.
(3) Storage for the microprogram (firmware).
(4) A register that holds the address of the microinstruction to be executed.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-3, R-2, S-1
(C) P-4, Q-3, R-1, S-2
(D) P-3, Q-4, R-1, S-2

124. What is the most significant advantage of microprogrammed control over hardwired control?
(A) Higher execution speed.
(B) Lower hardware cost for simple CPUs.
(C) Greater flexibility to modify or add new instructions.
(D) Reduced power consumption.

125. A microprogram is best defined as:
(A) A very small user program.
(B) A sequence of microinstructions that implements a single machine instruction.
(C) The hardware that generates control signals.
(D) The part of the operating system that handles interrupts.

126. Which of the following is an incorrect statement about address sequencing in a microprogrammed control unit?
(A) The next address can be the current address incremented by one.
(B) The next address can be determined by the opcode of the machine instruction.
(C) The next address is always fetched from the main memory.
(D) Conditional branching can be used to choose between alternative next addresses.

127. In a microprogrammed control unit, the size of the control memory is determined by:
(A) The number of instructions in the user program.
(B) The number of microinstructions and the width of each microinstruction.
(C) The size of the main memory.
(D) The number of CPU registers.

Click Here to Buy book on Amazon

128. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):

Assertion (A) : Microprogrammed control provides flexibility in designing a CPU's instruction set.

Reason (R) : The instruction set can be changed by modifying the microprogram in the control memory, without changing the hardware.

In the context of the two statements, which one of the following is correct?

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

129. The register that specifies the address of the microinstruction to be read from control memory is called the:

(A) Program Counter (PC)
(B) Instruction Register (IR)
(C) Control Address Register (CAR)
(D) Memory Address Register (MAR)

130. Given the following statements regarding microprogramming:

Statement I : Horizontal microprogramming allows for a higher degree of parallelism as multiple microoperations can be specified in one microinstruction.

Statement II : Vertical microprogramming uses a more encoded format, resulting in shorter microinstructions but requiring a decoder to generate control signals.

Which of the following is correct?

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

131. What is a key feature of a General Register Organization in a CPU?

(A) All operations occur between a single accumulator and memory.
(B) Any of the registers can be used as a source or destination for ALU operations.
(C) It requires a stack to perform any arithmetic operation.
(D) There is only one general-purpose register available to the programmer.

132. In a stack-organized computer, which operation involves removing an item from the top of the stack?

(A) PUSH
(B) POP
(C) LOAD
(D) STORE

133. Given the following statements regarding processor architectures:

Statement I : RISC architectures aim to execute simple instructions in a single clock cycle.

Statement II : CISC architectures have complex instructions that may take multiple clock cycles to execute.

Which of the following is correct?

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

134. Arrange the following actions into the correct chronological sequence for executing an instruction like `ADD R1, [X]`, which adds the content of memory location X to register R1:

(a) Decode the instruction.
(b) Fetch the instruction from memory.
(c) Fetch the operand from memory location X.
(d) Perform the addition and store the result in R1.

Choose the correct option:

(A) b, a, c, d
(B) a, b, c, d
(C) b, c, a, d
(D) b, a, d, c

135. Which of the following is a correct statement about instruction formats?

(A) The number of address fields in an instruction format depends on the internal organization of the CPU.
(B) A zero-address instruction format always uses two registers.
(C) All instruction formats must have the same fixed length.
(D) The opcode field is an optional part of an instruction format.

NODIA

136. A computer uses an addressing mode where the effective address (EA) is calculated as EA = A + (R), where A is an address part of the instruction and (R) is the content of a CPU register. What is this addressing mode called?
(A) Indexed Addressing
(B) Indirect Addressing
(C) Immediate Addressing
(D) Direct Addressing

137. Which of the following addressing modes does not require a memory reference to fetch an operand?
(A) Direct Addressing
(B) Indirect Addressing
(C) Register Indirect Addressing
(D) Immediate Addressing

138. Consider the following characteristics:
(a) Large number of general-purpose registers.
(b) Simple and limited instruction set.
(c) Instructions have variable lengths and complex formats.
(d) Emphasis on hardwired control unit.
Which of the above are considered typical characteristics of a RISC computer?
(A) (a), (c) and (d) only
(B) (b) and (c) only
(C) (a), (b) and (d) only
(D) (a), (b), (c) and (d)

139. What is the most appropriate purpose of having different addressing modes in a CPU?
(A) To increase the underlying complexity of the processor hardware design.
(B) To provide flexibility for accessing operands and managing data structures.
(C) To ensure that every instruction takes exactly the same number of clock cycles.
(D) To reduce the total number of general-purpose registers needed by the system.

140. In a stack organization, what does the Stack Pointer (SP) register hold?
(A) The data item at the top of the stack.
(B) The address of the top of the stack.
(C) The total number of items in the stack.
(D) The base address of the stack in memory.

141. Which of the following is an incorrect statement about CISC computer characteristics?
(A) They have a large and complex instruction set.
(B) Instructions often perform multiple low-level operations.
(C) They predominantly use a hardwired control unit for speed.
(D) They support a wide variety of addressing modes.

142. Match the Addressing Mode in List I with the way the Effective Address (EA) is determined in List II.
List I: (Addressing Mode)
(P) Immediate
(Q) Direct
(R) Register Indirect
(S) Relative
List II: (Effective Address)
(1) EA = Address part of instruction
(2) Operand is part of the instruction
(3) EA = Content of a CPU register
(4) EA = PC + Address part of instruction
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-1, R-3, S-4
(D) P-3, Q-1, R-2, S-4

143. RISC computers often include a large number of registers. What is the primary cause-and-effect relationship behind this design choice?
(A) Cause: To support complex addressing modes. Effect: Slower instruction execution.
(B) Cause: To minimize memory access operations. Effect: Faster program execution.
(C) Cause: To simplify the control unit. Effect: Increased need for memory accesses.
(D) Cause: To reduce the size of the instruction set. Effect: More complex compiler design.

144. An instruction format that does not have an address field is known as a:
(A) One-address instruction format
(B) Two-address instruction format
(C) Zero-address instruction format
(D) Three-address instruction format

145. Given the following statements regarding addressing modes:
Statement I : In Register Indirect addressing,

Downloaded from www.pdf.tube

the address field of the instruction specifies a register that contains the memory address of the operand.

Statement II : In Base Register addressing, the effective address is computed by adding a displacement to the content of a base register.

Which of the following is correct?

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

146. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):

Assertion (A) : Stack-organized CPUs are well-suited for evaluating arithmetic expressions written in Reverse Polish Notation (RPN).

Reason (R) : RPN expressions can be evaluated by pushing operands onto the stack and applying operators to the top elements without needing to manage temporary variables explicitly.

In the context of the two statements, which one of the following is correct?

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

147. What is the correct Reverse Polish Notation (RPN) for the infix expression (A + B) * (C + D)?

(A) A B + C D + *
(B) A B C D + * +
(C) * + A B + C D
(D) A B + * C D +

148. Identify the correct pair of instruction format and a typical operation.

(A) Three-address: ADD R1, R2 (R1 <- R1 + R2)
(B) Two-address: ADD R1, R2, R3 (R1 <- R2 + R3)
(C) PUSH A
(D) Zero-address: ADD (AC <- AC + M)

149. The Program Counter (PC) is a CPU register whose primary function is to:

(A) Store the instruction currently being executed.
(B) Hold the address of the next instruction to be fetched.
(C) Store the result of the last arithmetic operation.
(D) Hold the status of the CPU.

150. Consider the following fields of an instruction format:

(a) Opcode field
(b) Mode field for addressing
(c) Address or operand field
(d) Control memory address field
(e) A field for the program's author

Which of the above fields are commonly found?

(A) (a), (b), and (c) only
(B) (a), (c), and (d) only
(C) (a), (b), (c), and (d) only
(D) (a), (b), and (e) only

151. An instruction `ADD X` in a one-address machine typically implies the operation:

(A) X <- X + AC
(B) PC <- PC + X
(C) R1 <- R1 + M[X]
(D) AC <- AC + M[X]

152. Which of the following is an incorrect statement regarding the benefits of a General Register Organization?

(A) It reduces the number of memory accesses needed.
(B) It simplifies compiler design by providing more temporary storage options.
(C) It increases the length of instructions as registers need to be specified.
(D) It guarantees that all programs will run faster than on an accumulator-based machine.

153. If a stack is implemented to grow towards lower memory addresses, a PUSH operation would involve:

(A) Incrementing the stack pointer, then writing data.
(B) Writing data, then incrementing the stack pointer.
(C) Decrementing the stack pointer, then writing data.
(D) Writing data, then decrementing the stack pointer.

154. Which characteristic most clearly distinguishes a RISC computer from a CISC computer?

NODIA

(A) The number of I/O ports.
(B) The use of a memory hierarchy.
(C) The complexity of the instruction set and control unit.
(D) The ability to perform parallel processing.

155. Given the following statements regarding addressing modes:
Statement I : In Implied mode, the operands are specified implicitly in the definition of the instruction.
Statement II : In Immediate mode, the address field of the instruction contains the actual operand.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

156. What is an 'addressing mode'?
(A) The way a computer connects to a network.
(B) A method for specifying the location of the operands used by an instruction.
(C) The physical layout of memory chips on the motherboard.
(D) The protocol used for serial communication.

157. In a stack-organized computer, an instruction like 'ADD' with no address field would perform which operation?
(A) It would add the top two elements of the stack, pop them, and push the result.
(B) It would add a value from memory to the top of the stack.
(C) It would add the contents of two predefined registers.
(D) It would prompt the user for a number to add to the accumulator.

158. Which of the following statements about instruction types is correct?
(A) Data transfer instructions perform arithmetic and logic operations.
(B) Program control instructions are used to move data between memory and registers.
(C) Data manipulation instructions include arithmetic, logical, and shift operations.
(D) I/O instructions are a subset of data manipulation instructions.

159. What is a primary reason for the variable length of instructions in a CISC architecture?
(A) To simplify the decoding process.
(B) To support a wide variety of complex addressing modes and operations.
(C) To ensure every instruction fits into a single memory word.
(D) To make the control unit hardwired and fast.

160. A CPU's ALU, Control Unit, and Registers are collectively known as the:
(A) Memory Management Unit
(B) Central Processing Unit
(C) I/O Subsystem
(D) System Bus

161. What is the fundamental concept behind Parallel Processing?
(A) Executing programs sequentially but very quickly.
(B) Using multiple processing units to solve a problem concurrently.
(C) Storing data in a parallel format in memory.
(D) Transferring data over a parallel bus instead of a serial one.

162. The technique of decomposing a sequential process into sub-operations, with each sub-operation being executed in a special dedicated segment that operates concurrently with all other segments, is called:
(A) Multitasking
(B) Multiprogramming
(C) Pipelining
(D) Vector processing

163. Given the following statements regarding computer pipelines:
Statement I : An arithmetic pipeline is designed to speed up floating-point operations.
Statement II : An instruction pipeline is designed to overlap the fetch, decode, and execute phases of instructions.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

Click Here to Buy book on Amazon

164. Consider the following issues related to instruction pipelining:
(a) Resource conflicts (structural hazards).
(b) Data dependencies (data hazards).
(c) Branch difficulties (control hazards).
(d) Power failures.
Which of the above statement(s) is/are correct regarding disruptions to the smooth flow of instructions?
(A) (a) and (d) only
(B) (a), (b) and (c) only
(C) (b) and (d) only
(D) (a), (b), (c) and (d)

165. Which statement correctly describes Vector Processing?
(A) Processing a single piece of data at a time.
(B) Performing the same operation on multiple data elements simultaneously.
(C) A method for processing text and character strings.
(D) A software technique for managing vectors in a programming language.

166. If a 4-stage pipeline has stage delays of 10ns, 12ns, 8ns, and 10ns, what is the clock cycle time of the pipelined processor?
(A) 10ns
(B) 12ns
(C) 40ns
(D) 8ns

167. All of the following are forms of parallel processing as per Flynn's classification EXCEPT:
(A) SISD (Single Instruction, Single Data)
(B) SIMD (Single Instruction, Multiple Data)
(C) MIMD (Multiple Instruction, Multiple Data)
(D) SSSD (Single Storage, Single Data)

168. Match the Pipeline Hazard in List I with its Cause in List II.
List I: (Hazard)
(P) Structural Hazard
(Q) Data Hazard
(R) Control Hazard
List II: (Cause)
(1) An instruction depends on the result of a previous instruction that is not yet complete.
(2) The pipeline makes a decision about which instruction to fetch next before a branch is resolved.
(3) Two instructions require the same hardware resource at the same time.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3
(B) P-3, Q-1, R-2
(C) P-2, Q-3, R-1
(D) P-3, Q-2, R-1

169. What is the most appropriate advantage of instruction pipelining?
(A) It decreases the execution time required for a single machine instruction.
(B) It increases the instruction throughput of the processing unit.
(C) It reduces the overall logic complexity of the internal control unit.
(D) It eliminates the need for maintaining an internal program counter.

NODIA

170. An Array Processor is a type of processor that belongs to which category of Flynn's classification?
(A) SISD
(B) SIMD
(C) MISD
(D) MIMD

171. Which of the following is an incorrect statement about the operation of an instruction pipeline?
(A) It allows for the overlapping of instruction execution phases.
(B) The pipeline is divided into segments, each performing a part of the instruction cycle.
(C) The time taken to execute a single instruction from start to finish (latency) is significantly reduced.
(D) Branch instructions can cause the pipeline to be flushed and refilled.

172. A data hazard in a pipeline is caused by:
(A) Two instructions trying to access the same memory location simultaneously.
(B) An instruction needing the result of a preceding instruction that has not yet been computed.
(C) A branch instruction changing the normal flow of the program.
(D) The pipeline running out of instructions to execute.

173. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):

Assertion (A) : Instruction pipelining increases the throughput of a CPU.

Reason (R) : Pipelining allows multiple instructions to be in different stages of execution simultaneously.

In the context of the two statements, which one of the following is correct?

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

174. The concept of vector processing, where a single instruction operates on a collection of data, is characteristic of which architectural style?

(A) SISD
(B) MIMD
(C) SIMD
(D) MISD

175. Given the following statements regarding instructions:

Statement I : A vector instruction specifies an operation to be performed on an entire vector of data.

Statement II : A scalar instruction operates on a single data item.

Which of the following is correct?

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

176. Consider the following characteristics of vector processing:

(a) Instructions operate on arrays or vectors of data.
(b) It is highly effective for scientific and engineering computations.
(c) It typically involves deeply pipelined functional units.
(d) Each element of the vector is processed by a separate processor in MIMD fashion.
(e) It is primarily used for text processing and database management.

Which of the above statement(s) is/are correct?

(A) (d) and (e) only
(B) (a), (b), and (c) only
(C) (a), (c), and (e) only
(D) (a), (b), (c) and (d)

177. A space-time diagram for a pipeline shows:

(A) The physical layout of the pipeline stages.
(B) The memory usage of the pipeline over time.
(C) The utilization of each pipeline stage over a sequence of clock cycles.
(D) The power consumption of the pipeline.

178. Identify the correct pair of a processor type and its key characteristic.

(A) Scalar Processor - Executes one instruction on multiple data streams.
(B) Vector Processor - Executes multiple instructions on a single data stream.
(C) Array Processor - A synchronized set of processors executing the same instruction on different data.
(D) Multiprocessor - A single CPU executing multiple programs.

179. The purpose of an Arithmetic Pipeline is to:

(A) Overlap the fetch-decode-execute cycle of different instructions.
(B) Speed up the execution of complex arithmetic operations like floating-point multiplication.
(C) Process arrays of data simultaneously.
(D) Handle I/O operations in a pipelined fashion.

180. The technique of branch prediction is used in instruction pipelines to:

(A) Reduce the penalty incurred from control hazards.
(B) Resolve data dependencies between instructions.
(C) Avoid structural hazards.
(D) Increase the accuracy of arithmetic calculations.

NODIA

181. What are peripheral devices in a computer system?

(A) The core components of the CPU like the ALU and control unit.
(B) External devices such as keyboards, monitors, and printers that communicate

with the computer.
(C) Internal memory modules like RAM and cache.
(D) Software programs that run on the computer.

182. What is the primary function of an Input-Output Interface?
(A) To resolve differences between the CPU and peripherals, and supervise data transfer.
(B) To execute application programs and manage high-level software processes.
(C) To store data permanently within the system's non-volatile storage units.
(D) To increase the base clock speed of the central processing unit's clock.

183. Given the following statements regarding asynchronous data transfer:
Statement I : The strobe control method uses a single control line to time each transfer.
Statement II : The handshaking method uses two control lines to provide a confirmation signal, making it more reliable.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

184. Arrange the sequence of events for a typical DMA block transfer from an I/O device to memory in the correct order:
(a) DMA controller takes control of the bus from the CPU.
(b) CPU initializes the DMA controller with the memory address, word count, and I/O device address.
(c) DMA controller transfers the data block directly to memory.
(d) DMA controller interrupts the CPU upon completion of the transfer.
Choose the correct option:
(A) b, a, c, d
(B) a, b, c, d
(C) b, c, a, d
(D) a, c, d, b

185. Which of the following is a correct statement about the modes of data transfer?
(A) Programmed I/O is the most efficient mode for transferring large blocks of data.
(B) Interrupt-initiated I/O allows the CPU to execute other programs while waiting for I/O.
(C) Direct Memory Access (DMA) requires constant intervention from the CPU for every byte transferred.
(D) All I/O transfers must go through the accumulator.

186. In a system with multiple devices connected to a single interrupt request line, what is the purpose of a priority interrupt system?
(A) To ensure all devices are serviced in a first-come, first-served order.
(B) To allow a higher-priority device to interrupt the service of a lower-priority device.
(C) To disable all interrupts from occurring.
(D) To allow only one device to ever generate an interrupt.

187. All of the following are standard modes of data transfer between the CPU/memory and peripherals EXCEPT:
(A) Programmed I/O
(B) Interrupt-initiated I/O
(C) Direct Memory Access (DMA)
(D) Pipelined I/O

188. Consider the following statements about Direct Memory Access (DMA):
(a) The DMA controller manages data transfers without CPU intervention for each word.
(b) It is particularly efficient for high-speed, large-volume data transfers.
(c) The CPU is completely idle during the entire DMA process.
(d) It requires special registers in the DMA controller for address, word count, and control.
Which of the above statement(s) is/are correct?
(A) (a) and (c) only
(B) (b), (c) and (d) only
(C) (a), (b) and (d) only
(D) (a), (b), (c) and (d)

189. What is the most appropriate reason for using an I/O interface between the CPU and a peripheral?
(A) To convert digital signals to analog for processing by the central unit.
(B) To increase the total storage capacity of the primary main memory unit.

Downloaded from www.pdf.tube

(C) To resolve the differences in speed, data format, and electrical characteristics.
(D) To provide a specific programming language used only for the peripheral.

190. A DMA controller is a specialized processor that:
(A) Executes user programs.
(B) Manages data transfers between memory and I/O devices.
(C) Controls the cache memory.
(D) Performs floating-point arithmetic.

191. Which of the following statements about priority interrupt schemes is incorrect?
(A) Polling involves the CPU querying each device to identify the source of the interrupt.
(B) In a daisy-chaining scheme, the priority is determined by the physical position of the device on the chain.
(C) A parallel priority scheme uses a register where each bit corresponds to an interrupt request.
(D) Polling is the fastest method for identifying the highest-priority interrupt.

192. Match the I/O Terminology in List I with its Definition in List II.
List I: (Term)
(P) Handshaking
(Q) Memory-mapped I/O
(R) DMA
(S) Polling
List II: (Definition)
(1) An I/O scheme where memory and I/O devices share the same address space.
(2) A method where the CPU repeatedly checks the status of a device.
(3) A data transfer scheme involving an exchange of control signals for synchronization.
(4) A mechanism for data transfer without direct CPU involvement for each byte.
Choose the correct option from those given below:
(A) P-3, Q-1, R-4, S-2
(B) P-1, Q-3, R-4, S-2
(C) P-3, Q-2, R-4, S-1
(D) P-2, Q-1, R-4, S-3

193. The mode where the DMA controller takes control of the system bus from the CPU for one or more clock cycles to transfer data is known as:
(A) Burst mode
(B) Cycle stealing
(C) Transparent mode
(D) Interrupt mode

194. Which communication method transmits data one bit at a time over a single communication line?
(A) Parallel communication
(B) Serial communication
(C) DMA transfer
(D) Bus transfer

195. Given the following statements regarding I/O configurations:
Statement I : In isolated I/O, the CPU has distinct input and output instructions, and the I/O port addresses are separate from memory addresses.
Statement II : In memory-mapped I/O, I/O devices are treated like memory locations, and the same instructions can be used to access both memory and I/O devices.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

196. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R):
Assertion (A) : Direct Memory Access (DMA) is used for high-speed data transfer between peripherals and main memory.
Reason (R) : DMA transfers data directly without involving the CPU in the transfer of each word, thus freeing the CPU to perform other tasks.
In the context of the two statements, which one of the following is correct?
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

197. A real-time system needs to respond very quickly to data from a high-speed sensor. Which I/O transfer mode would be the most suitable?
(A) Programmed I/O
(B) Interrupt-driven I/O
(C) Direct Memory Access (DMA)
(D) Polling-based I/O

198. Identify the correct pair of a data transfer mode and its primary characteristic.
(A) Programmed I/O - CPU is freed to do other tasks.
(B) Interrupt-initiated I/O - CPU must wait in a loop.
(C) DMA - Reduces CPU overhead for bulk data transfer.
(D) Serial Transfer - Best for slow, character-by-character transfers.

199. What is 'handshaking' in the context of asynchronous data transfer?
(A) The physical connection of a device to the bus.
(B) An exchange of control signals between the source and destination units to ensure synchronized data transfer.
(C) The process of the CPU checking the status of an I/O device.
(D) The initial setup of a DMA transfer.

200. Consider the following functions related to an I/O interface:
(a) To provide a buffer for storing data temporarily.
(b) To convert data from serial to parallel format or vice-versa.
(c) To generate the control signals for the I/O transfer.
(d) To decode the address of the peripheral device.
(e) To execute the main application logic.
Which of the above are typical functions of an I/O interface?
(A) (e) only
(B) (a), (b), and (c) only
(C) (a), (b), (c), and (d) only
(D) (a) and (d) only

201. The 'daisy-chaining' method is used for:
(A) Connecting multiple memory modules.
(B) Implementing a priority interrupt system.
(C) Transferring data in serial communication.
(D) Organizing a cache memory.

202. Which of the following statements about asynchronous data transfer is incorrect?
(A) It does not use a common clock between the sender and receiver.
(B) The handshaking protocol is a method for asynchronous transfer.
(C) It is suitable for devices with widely different speeds.
(D) It is always faster than synchronous data transfer.

203. A serial communication interface is configured with a baud rate of 9600, 8 data bits, 1 stop bit, and no parity. Approximately how many characters per second can be transmitted?
(A) 960
(B) 9600
(C) 1200
(D) 800

204. What is the main principle of a priority interrupt system?
(A) To service all interrupts in the order they arrive.
(B) To establish a hierarchy for handling simultaneous interrupt requests from multiple devices.
(C) To disable all interrupts while the CPU is busy.
(D) To handle interrupts using software polling only.

205. In serial communication, what is the purpose of the start and stop bits?
(A) To carry the actual data.
(B) To indicate the beginning and end of a character for synchronization.
(C) To check for transmission errors.
(D) To specify the address of the destination device.

206. Which type of memory constitutes the directly accessible, primary storage for the CPU?
(A) Auxiliary Memory
(B) Main Memory
(C) Cache Memory
(D) Associative Memory

Click Here to Buy book on Amazon

207. What is a key characteristic of auxiliary memory, such as a hard disk or SSD?
(A) It is volatile.
(B) It has a faster access time than main memory.
(C) It is non-volatile and has a large storage capacity.
(D) It is directly accessible by the CPU for instruction execution.

208. Given the following statements regarding associative memory:
Statement I : Associative memory is accessed by its content rather than by its address.
Statement II : Cache memory is a type of associative memory.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

209. Consider the following mapping techniques for cache memory:
(a) Direct mapping
(b) Associative mapping
(c) Set-associative mapping
(d) Sequential mapping
Which of the above are common mapping techniques?
(A) (a) and (d) only
(B) (a), (b) and (c) only
(C) (b) and (d) only
(D) (a), (b), (c) and (d)

210. Which of the following is a correct statement about the concept of virtual memory?
(A) It increases the total amount of physical RAM in a computer.
(B) It allows the execution of programs that may be larger than the physical main memory.
(C) It is a very fast memory that is located inside the CPU.
(D) It is a form of non-volatile memory used for long-term storage.

211. A cache system has a hit ratio of 0.9. The cache access time is 10 ns, and the main memory access time is 100 ns. What is the average memory access time?
(A) 19 ns
(B) 20 ns
(C) 90 ns
(D) 55 ns

212. A typical memory hierarchy is organized based on speed, cost, and capacity. Which of the following is NOT a property used in this organization?
(A) Access time
(B) Cost per bit
(C) Physical location on the motherboard
(D) Volatility

NODIA

213. Arrange the following memory types in the correct order of decreasing access speed (fastest first):
(a) CPU Registers
(b) Cache Memory
(c) Main Memory (DRAM)
(d) Magnetic Disk (Auxiliary Memory)
Choose the correct option:
(A) a, c, b, d
(B) d, c, b, a
(C) a, b, c, d
(D) b, a, c, d

214. What is the most appropriate purpose of cache memory?
(A) To provide long-term, non-volatile storage for large data files and system documentation.
(B) To act as a buffer between the CPU and main memory, holding frequently accessed data.
(C) To extend the total addressable memory space beyond the physical capacity of the system.
(D) To store the initial boot-up sequence and firmware settings of the computer system.

215. What is the role of Memory Management Hardware (MMU)?
(A) To perform arithmetic and logic operations.
(B) To translate logical addresses generated by the CPU into physical addresses in main memory.
(C) To control the flow of data on the system bus.
(D) To store the operating system's kernel code.

216. Which of the following statements about virtual memory implementation is incorrect?
(A) It typically uses a hard disk as part of the memory hierarchy.
(B) A page table is used to map virtual pages to physical frames.
(C) A page fault occurs when a referenced page is not in main memory.
(D) The entire program must be loaded from disk into memory before it can begin execution.

217. The principle of locality of reference is fundamental to the effectiveness of cache memory. This principle implies that:
(A) All memory references are to the same location.
(B) The CPU accesses memory in a completely random pattern.
(C) The CPU tends to access memory locations that are near other recently accessed locations or are accessed repeatedly.
(D) Data and instructions must be located in adjacent memory blocks.

218. Consider the following properties of memory hierarchy:
(a) As we go down the hierarchy (from registers to tape), the cost per bit decreases.
(b) As we go down the hierarchy, the access time increases.
(c) As we go down the hierarchy, the capacity increases.
(d) As we go up the hierarchy (from tape to registers), the processor accesses the memory more frequently.
(e) All levels of the hierarchy are volatile.
Which of the above statement(s) is/are correct?
(A) (a), (b) and (e) only
(B) (c), (d) and (e) only
(C) (a), (b), (c) and (d) only
(D) (a), (c) and (d) only

219. The 'write-through' and 'write-back' policies are associated with which memory component?
(A) Virtual Memory
(B) Auxiliary Memory
(C) Cache Memory
(D) Associative Memory

220. Given the following statements regarding semiconductor memories:
Statement I : Static RAM (SRAM) is typically used for cache memory because it is faster than Dynamic RAM (DRAM).
Statement II : Dynamic RAM (DRAM) requires periodic refreshing to maintain its stored data.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

221. In a set-associative mapped cache with 2 blocks per set, a main memory address is divided into three fields: Tag, Set, and Word. To which set would the main memory block 13 map in a cache with 8 sets?
(A) Set 6
(B) Set 1
(C) Set 5
(D) Set 3

222. Identify the correct pair matching a memory type with its typical access time characteristic.
(A) CPU Registers - Milliseconds
(B) Main Memory (DRAM) - Microseconds
(C) Hard Disk Drive - Picoseconds
(D) Cache (SRAM) - Nanoseconds

223. What is a 'page table' in a virtual memory system?
(A) A hardware device for page replacement.
(B) A data structure in memory that maps virtual page numbers to physical frame numbers.
(C) A table of contents for a program.
(D) A part of the cache memory that stores pages.

224. The search operation in an associative memory is a:
(A) Sequential search by address.
(B) Parallel search by content.
(C) Binary search by value.
(D) Direct access by index.

225. Which of the following is an incorrect statement regarding memory management?
(A) The LRU (Least Recently Used) algorithm is a common page replacement policy.

Downloaded from www.pdf.tube

(B) A page fault is a software interrupt handled by the OS.
(C) A Translation Lookaside Buffer (TLB) is a cache for the page table.
(D) Virtual memory is always faster to access than physical main memory.

226. In a paged virtual memory system, the CPU generates a logical address consisting of a page number and an offset. The Memory Management Unit (MMU) uses the page number to:
(A) Index into the page table to find the corresponding frame number.
(B) Access an I/O device.
(C) Calculate the physical address directly.
(D) Determine the size of the page.

227. Thrashing in a virtual memory system refers to a situation where:
(A) The hit ratio of the cache is nearly 100%.
(B) The CPU is executing instructions at its maximum speed.
(C) A program is stuck in an infinite loop.
(D) The system is spending an excessive amount of time swapping pages between main memory and auxiliary memory.

228. Which cache replacement policy removes the block that has been in the cache for the longest time without being referenced?
(A) Least Recently Used (LRU)
(B) First-In, First-Out (FIFO)
(C) Random Replacement
(D) Last-In, First-Out (LIFO)

229. What is a 'cache hit'?
(A) An event where the CPU finds the required data in main memory.
(B) An event where the CPU finds the required data in the cache.
(C) An event where the data is not found in either cache or main memory.
(D) An event where a cache block is replaced.

230. Which of the following statements about auxiliary memory is correct?
(A) It is typically volatile, losing its contents when power is turned off.
(B) Magnetic tapes and optical disks are examples of auxiliary memory devices.
(C) It is more expensive per bit than main memory.
(D) It is accessed directly by the CPU for program execution.

231. Which of the following is a primary characteristic of a multiprocessor system?
(A) It contains a single CPU with multiple cores.
(B) It consists of multiple independent computers connected by a network.
(C) It contains two or more CPUs that share access to a common memory and peripherals.
(D) It is a theoretical model that cannot be built in practice.

232. Which of the following is an example of an interconnection structure for multiprocessors?
(A) A crossbar switch
(B) A stack
(C) A page table
(D) An instruction register

233. Given the following statements regarding multiprocessors:
Statement I : Tightly coupled multiprocessors share a single main memory and are controlled by one operating system.
Statement II : Loosely coupled multiprocessors, or distributed systems, communicate through message passing over a network.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

234. Consider the following mechanisms for interprocessor arbitration:
(a) Serial arbitration procedure (daisy-chaining).
(b) Parallel arbitration logic.
(c) Dynamic arbitration algorithms like LRU or FIFO.
(d) Using a semaphore.
Which of the above mechanism(s) is/are correct?
(A) (a) and (d) only
(B) (a), (b) and (c) only
(C) (d) only
(D) (a), (b), (c) and (d)

NODIA

235. Which statement correctly describes the need for interprocessor communication and synchronization?
(A) It is only required in single-processor systems.
(B) It is necessary to coordinate activities and ensure correct data sharing among processors in a multiprocessor system.
(C) It is a hardware feature that replaces the operating system.
(D) It is primarily used to manage virtual memory.

236. A multiprocessor system has two CPUs, each with its own local cache. CPU1 reads memory location X. Later, CPU2 writes a new value to location X. If CPU1 reads location X again, it may get the old, stale data from its cache. This problem is known as:
(A) The cache coherence problem.
(B) The bus arbitration problem.
(C) The thrashing problem.
(D) The deadlock problem.

237. All of the following are common interconnection structures for multiprocessors EXCEPT:
(A) Time-shared common bus
(B) Multiport memory
(C) Crossbar switch
(D) Single-port register file

238. Which of the following best defines a multicore processor?
(A) A system with multiple separate physical processor chips installed on the motherboard.
(B) A chip containing two or more independent processing units known as cores.
(C) A specialized processor that strictly utilizes vector processing for its data operations.
(D) A computer system configured with multiple hard drives for data storage redundancy.

239. The 'Cache Coherence' problem in multiprocessors refers to the issue of:
(A) Ensuring that all caches have the most recent and correct value of a shared data item.
(B) Ensuring that the cache hit rate is above a certain threshold.
(C) Ensuring that the cache and main memory sizes are proportional.
(D) Ensuring that all processors have a cache.

240. Which of the following statements about interprocessor arbitration is incorrect?
(A) It is the process of resolving requests from multiple processors for a shared resource, like a bus.
(B) The daisy-chaining method is a form of parallel arbitration.
(C) A static priority scheme assigns a fixed priority to each processor.
(D) A rotating priority (round-robin) scheme gives all processors an equal chance to access the resource.

241. The need for synchronization primitives like semaphores or mutexes in multiprocessor systems is primarily caused by:
(A) The need to prevent multiple processors from simultaneously modifying shared data, which could lead to inconsistent results.
(B) The need to increase the speed of individual processors.
(C) The requirement to connect processors using a bus.
(D) The high cost of cache memory.

242. Consider the following characteristics of multiprocessor systems:
(a) They aim to increase throughput and computational power.
(b) They can improve reliability through graceful degradation.
(c) They require mechanisms for communication and synchronization between processors.
(d) They always use a single-bus interconnection structure.
(e) They are classified as SISD systems under Flynn's taxonomy.
Which of the above statement(s) is/are correct?
(A) (d) and (e) only
(B) (a), (b), and (c) only
(C) (a), (c), and (e) only
(D) (a), (b), (c) and (d)

243. The 'snoopy protocol' or 'bus snooping' is a technique used to:
(A) Monitor network traffic for security threats.
(B) Maintain cache coherence in bus-based

Click Here to Buy book on Amazon

multiprocessor systems.
(C) Arbitrate access to the system bus.
(D) Debug programs running on a multiprocessor.

244. Given the following statements regarding multiprocessor architectures:
Statement I : In a Uniform Memory Access (UMA) architecture, all processors have equal access time to any part of the shared memory.
Statement II : In a Non-Uniform Memory Access (NUMA) architecture, the access time to a memory location depends on which processor is making the request and the location of the memory.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

245. Identify the correct pair of an interconnection network and its topology.
(A) Crossbar Switch - A single shared path for all components.
(B) Common Bus - A mesh-like structure with dedicated paths.
(C) Hypercube - A multidimensional binary cube structure.
(D) Multiport Memory - A daisy-chained ring structure.

246. What is 'Interprocessor Arbitration'?
(A) The mechanism for deciding which processor gets access to a shared resource when there are multiple simultaneous requests.
(B) The process of processors negotiating a common task.
(C) The communication protocol used between processors.
(D) The method for keeping processor caches coherent.

247. Which of the following is the most significant challenge in designing and programming large-scale multiprocessor systems?
(A) Providing adequate power and cooling resources for the hardware modules.
(B) Finding enough physical space for the hardware in the data center.
(C) Managingcommunication,synchronization, and data consistency among the many processors.
(D) Designing a user-friendly command-line interface for the system administrator.

248. An interconnection structure that provides a separate path from each memory module to each processor, eliminating bus contention but with high hardware cost, is the:
(A) Common bus
(B) Crossbar switch
(C) Hypercube network
(D) Ring network

249. Which of the following is a key advantage of multicore processors over single-core processors?
(A) They provide increased performance for parallel tasks within a smaller power and thermal envelope.
(B) They can run a single-threaded program significantly faster than any single-core model.
(C) They simplify the operating system's internal task scheduling and resource allocation logic.
(D) They do not require any internal cache memory to operate effectively at high speeds.

250. In a multiprocessor environment, a semaphore is a variable used for:
(A) Counting the number of processors in the system.
(B) Controlling access to a shared resource by multiple processes to achieve synchronization.
(C) Storing the address of the next instruction to execute.
(D) Broadcasting messages to all processors.

Direction : Please read the following source carefully.

The foundational layer of digital computers is constructed from digital logic circuits, which are themselves built from elementary components known as logic gates. These gates—such as AND, OR, and NOT—perform basic logical operations on binary inputs. The behavior and interconnection of these gates are mathematically

Downloaded from www.pdf.tube

described by Boolean Algebra, a system of logic that provides the rules and theorems necessary for the analysis and synthesis of digital circuits. For instance, De Morgan's theorems are instrumental in simplifying complex logical expressions. To further aid in the design of efficient circuits, Map Simplification techniques, most notably the Karnaugh Map (K-map), offer a graphical method for minimizing Boolean expressions. This simplification is crucial as it reduces the number of required logic gates, leading to lower cost, power consumption, and propagation delay. These fundamental elements are combined to create two primary categories of circuits: combinational circuits, where the output is a pure function of the present inputs (e.g., decoders, multiplexers), and sequential circuits, where the output depends on both present inputs and the past sequence of inputs, requiring memory elements like flip-flops. The entire framework, from basic gates to complex integrated circuits, relies on the principles of Boolean logic and its systematic simplification.

Based on the above source, attempt the following five questions.

251. According to the passage, what is the primary purpose of applying Boolean Algebra and Map Simplification techniques in digital circuit design?
(A) To increase the operational speed of logic gates exclusively.
(B) To reduce the complexity and resource requirements of the circuit.
(C) To allow for the use of both analog and digital signals.
(D) To enable the representation of data in a ternary number system.

252. Which of the following statements can be inferred about the relationship between combinational and sequential circuits?
(A) Combinational circuits lack memory, whereas sequential circuits incorporate it.
(B) Sequential circuits are a subtype of combinational circuits that include memory.
(C) Both circuit types produce outputs dependent solely on their current inputs.
(D) Combinational circuits use flip-flops, while sequential circuits use multiplexers.

253. The passage suggests that the use of a Karnaugh Map is a method for:
(A) performing arithmetic operations in binary.
(B) visualizing the physical layout of an integrated circuit.
(C) systematically minimizing a logical expression.
(D) converting digital signals to analog signals.

254. Based on the description, a decoder, which generates a specific output for a given binary input code, would be classified as what type of circuit?
(A) A sequential circuit, as it must remember previous codes.
(B) A memory unit, as it stores the output for future use.
(C) A logic gate, as it performs a single Boolean operation.
(D) A combinational circuit, as its output depends only on the present input.

255. What fundamental concept underpins the entire process of designing and simplifying digital circuits, from individual gates to integrated circuits?
(A) The principles of floating-point representation.
(B) The architecture of register transfer language.
(C) The rules and theorems of Boolean logic.
(D) The mechanics of asynchronous data transfer.

NODIA

Direction : Please read the following source carefully.

Data representation within a digital computer is fundamental to its operation, dictating how information is stored and manipulated. Digital systems operate on binary data, but must represent various data types, including numbers, text, and instructions. Number systems, particularly binary, octal, and hexadecimal, are used, and the ability to perform conversion between these systems is a crucial skill. For representing signed integers, complement arithmetic is widely employed. The 2's complement representation is particularly prevalent because it simplifies subtraction into an addition operation and provides a unique representation for zero, unlike the 1's complement system. Numeric data can

be handled through two primary formats: fixed-point and floating-point representation. Fixed-point representation assumes a fixed position for the radix point, making it suitable for applications where the range of numbers is limited and computation must be fast, such as in digital signal processing. Conversely, floating-point representation, which uses a mantissa and an exponent, allows for a much wider dynamic range of values and is essential for scientific and engineering calculations. To ensure data integrity during storage and transmission, error detection codes, such as parity bits, are incorporated into the data representation to detect single-bit errors.

Based on the above source, attempt the following five questions.

256. What is the primary advantage of using 2's complement representation for signed integers as described in the passage?
(A) It allows for the representation of a larger range of positive numbers.
(B) It simplifies the hardware for subtraction by converting it to addition.
(C) It is easier for humans to read and write than other systems.
(D) It provides multiple representations for the number zero.

257. A scientific application requires calculations involving very large and very small numbers with high precision. Based on the passage, which data representation format would be most suitable?
(A) Fixed-point representation, due to its computational speed.
(B) 1's complement representation, due to its handling of negative values.
(C) An error detection code, due to its focus on data integrity.
(D) Floating-point representation, due to its wide dynamic range.

258. The passage mentions error detection codes. What can be inferred about their function?
(A) They ensure the integrity of data by flagging potential corruption.
(B) They correct any errors found in the data automatically.
(C) They convert data between different number systems.
(D) They compress data to save storage space.

259. Which statement accurately contrasts fixed-point and floating-point representations based on the provided text?
(A) Fixed-point is ideal for a wide range of values, while floating-point is for a limited range.
(B) Floating-point is simpler to implement in hardware than fixed-point.
(C) Fixed-point has an implied radix point, while floating-point uses an explicit exponent.
(D) Floating-point is primarily for integer arithmetic, while fixed-point is for real numbers.

260. The need for conversion between binary, octal, and hexadecimal number systems arises because:
(A) they offer different ways to represent the same underlying binary data.
(B) computer hardware can only perform floating-point arithmetic.
(C) these systems are required for implementing error detection codes.
(D) only hexadecimal can represent signed numbers using complements.

Direction : Please read the following source carefully.

The foundational principle of a modern digital computer is the stored program organization, where both program instructions and data are stored in the same main memory. This architecture is central to the Basic Computer Organization and Design. The operation of the computer is driven by a sequence of actions known as the instruction cycle. Each instruction cycle typically consists of several phases, beginning with fetching an instruction from memory. The address of the next instruction is held in a special Computer Register called the Program Counter (PC). Once fetched, the instruction is placed in the Instruction Register (IR), where its opcode is decoded by the Timing and Control unit. This unit generates the necessary control signals to execute the instruction. Execution may involve processing data stored in other registers, such as an accumulator, or it may require fetching an operand from memory, as specified by memory-reference instructions. After execution, the cycle

Click Here to Buy book on Amazon

repeats. This orderly progression can be altered by events like an interrupt, an external or internal signal that causes the CPU to suspend its current task and attend to a specific event, such as an input-output request. The entire process demonstrates a tightly synchronized interplay between computer registers, memory, and the control unit to execute a program.

Based on the above source, attempt the following five questions.

261. According to the stored program organization concept, where are program instructions located?
(A) In a separate, dedicated control memory.
(B) Permanently within the Timing and Control unit.
(C) Exclusively inside the Computer Registers.
(D) In the same memory as the data they operate on.

262. What is the correct sequence of events at the beginning of an instruction cycle as described in the passage?
(A) Decode instruction from IR, fetch instruction from memory, place instruction in PC.
(B) Fetch instruction using address in PC, place instruction in IR, decode instruction.
(C) Generate control signals, fetch operand from memory, execute the instruction.
(D) Place instruction in IR, increment PC, fetch instruction from memory.

263. The passage suggests that the primary role of the Timing and Control unit is to:
(A) perform arithmetic and logical operations on data.
(B) store the results of instruction execution temporarily.
(C) generate the signals required to direct the steps of the instruction cycle.
(D) hold the memory address of the next instruction to be fetched.

264. An interrupt is described as an event that causes the CPU to:
(A) permanently halt the execution of the current program.
(B) ignore all subsequent memory-reference instructions.
(C) immediately fetch the next instruction from the Program Counter.
(D) pause the current sequence of operations to handle a specific request.

265. Based on the description, what is the function of the Instruction Register (IR)?
(A) To hold the instruction that is currently being decoded and executed.
(B) To hold the address of the memory location being accessed.
(C) To accumulate the results of arithmetic operations.
(D) To count the number of instructions executed in a program loop.

NODIA

Direction : Please read the following source carefully.

The design of a computer's control unit can be approached in two distinct ways: hardwired and microprogrammed. A microprogrammed control unit is characterized by its use of a special memory unit called Control Memory. This memory stores a set of microprograms, which consist of sequences of microinstructions. Each machine language instruction initiates a corresponding microprogram whose microinstructions generate the elementary control signals required for execution. The process of selecting the next microinstruction is known as address sequencing. Various techniques for address sequencing allow for conditional branching and subroutines within the microcode, providing significant flexibility. For instance, the address of the next microinstruction can be determined by the opcode of the machine instruction, the status flags of the CPU, or a field within the current microinstruction itself. This design approach contrasts with a hardwired control unit, which uses a fixed network of logic gates to generate control signals. While a hardwired unit may be faster, the design of a microprogrammed control unit is more systematic and flexible, making it easier to implement complex instruction sets (as in CISC architectures) and to modify or correct the control logic by simply updating the microprogram in the control memory.

Based on the above source, attempt the following five questions.

266. What is the fundamental component that distinguishes a microprogrammed control unit from a hardwired one?
(A) The use of sequential logic circuits.
(B) The presence of a Control Memory.
(C) The ability to handle interrupts.
(D) The connection to a system bus.

267. Within a microprogrammed control unit, address sequencing refers to the process of:
(A) determining the memory location of the next machine instruction.
(B) selecting the address of the next microinstruction to be executed.
(C) calculating the address of an operand in main memory.
(D) translating a logical memory address to a physical address.

268. According to the passage, a significant advantage of using a microprogrammed control unit is:
(A) the relative ease of implementing complex instructions and modifying control logic.
(B) its inherently faster operational speed compared to hardwired units.
(C) its lower manufacturing cost due to the absence of memory.
(D) its exclusive reliance on combinational logic circuits for signal generation.

269. The passage implies that the control signals in a microprogrammed system are generated by:
(A) a fixed configuration of AND, OR, and NOT gates.
(B) direct hardware connections from the instruction register.
(C) the execution of a sequence of microinstructions.
(D) the computer's arithmetic and logic unit (ALU).

270. What is a microprogram in the context of a microprogrammed control unit?
(A) A high-level language program written by a user.
(B) A sequence of machine instructions for a specific task.
(C) An operating system routine for managing peripheral devices.
(D) A set of microinstructions stored in Control Memory to implement one machine instruction.

Direction : Please read the following source carefully.

The Central Processing Unit (CPU) is the brain of the computer, and its internal architecture significantly impacts performance. A General Register Organization is a common design where the CPU contains a set of general-purpose registers that can be used for various functions, such as storing operands or addresses. This is more flexible than an accumulator-based design, which relies on a single dedicated register for arithmetic results. An alternative is a Stack Organization, which manages data using a last-in, first-out (LIFO) stack. Operations are performed on the top elements of the stack, simplifying instruction formats as operands are implicitly at the top of the stack. The structure of machine instructions themselves, known as Instruction Formats, is critical. Formats define the layout of the bits in an instruction, specifying fields for the opcode, operand addresses, and the addressing mode. Addressing Modes are rules that determine how the operand of an instruction is specified. They provide versatility in accessing data, with modes ranging from immediate (operand is in the instruction), direct (address is in the instruction), and indirect (instruction contains an address pointing to the operand's address) to more complex modes like indexed and relative addressing. The choice of organization, instruction formats, and addressing modes collectively defines the CPU's programming model and efficiency.

Based on the above source, attempt the following five questions.

271. What is the primary advantage of a General Register Organization over an accumulator-based design?
(A) It requires fewer bits to represent an instruction format.
(B) It simplifies the control unit design by eliminating addressing modes.
(C) It offers greater flexibility by allowing multiple registers for operands and results.
(D) It relies exclusively on a stack for all data manipulations.

Downloaded from www.pdf.tube

272. In a Stack Organization CPU, where are the operands for an arithmetic instruction typically located?
(A) In general-purpose registers specified by the instruction.
(B) Implicitly at the top of the stack structure.
(C) At an address in main memory provided in the instruction.
(D) Within the instruction itself as an immediate value.

273. The term Addressing Modes refers to the rules that:
(A) define the total number of instructions a CPU can execute.
(B) manage the LIFO structure of a stack.
(C) decode the opcode field of an instruction.
(D) specify how the location of an operand is determined.

274. If an instruction uses an addressing mode where the instruction itself contains the operand's actual value, which mode is being used?
(A) Immediate addressing
(B) Indirect addressing
(C) Direct addressing
(D) Indexed addressing

275. What is the main purpose of defining different Instruction Formats?
(A) To ensure all instructions have the same length and execution time.
(B) To manage the flow of data between the CPU and peripheral devices.
(C) To structure the bits of an instruction to represent opcode, operands, and mode.
(D) To organize the memory hierarchy into cache, main, and auxiliary levels.

Direction : Please read the following source carefully.

To enhance computational speed, modern computers employ techniques for parallel processing, where multiple operations are performed concurrently. Pipelining is a particularly effective implementation of this concept. Instead of processing instructions sequentially one after another, an instruction pipeline overlaps the execution of multiple instructions. The instruction cycle is broken down into a series of stages, such as instruction fetch (IF), instruction decode (ID), operand fetch (OF), instruction execution (IE), and result write-back (WB). As one instruction moves from the fetch to the decode stage, the pipeline fetches the next instruction. In an ideal scenario, a new instruction can complete in every clock cycle, significantly increasing instruction throughput. This same principle can be applied to complex arithmetic operations in an arithmetic pipeline, where tasks like floating-point addition are broken into sub-operations (e.g., compare exponents, align mantissas, add mantissas, normalize result) that are pipelined. Another form of parallel computing involves vector processing, which operates on entire vectors of data (one-dimensional arrays) simultaneously using specialized hardware. This is highly efficient for scientific and engineering applications with a high degree of data parallelism. Array processors, which consist of multiple processing units under a common control, are a hardware implementation that excels at vector processing tasks.

Based on the above source, attempt the following five questions.

276. What is the primary goal of using an instruction pipeline?
(A) To increase the overall instruction throughput of the processor.
(B) To decrease the execution time of a single, individual instruction.
(C) To reduce the amount of memory required to store a program.
(D) To simplify the design of the computer's control unit.

277. The concept of pipelining is an example of what broader computational strategy?
(A) Asynchronous data transfer
(B) Virtual memory management
(C) Microprogrammed control
(D) Parallel processing

278. According to the passage, an arithmetic pipeline is specifically designed to:
(A) overlap the stages of fetching and decoding instructions.
(B) accelerate complex numerical calculations by breaking them into stages.
(C) manage data transfers between the CPU and I/O devices.
(D) process entire arrays of data using multiple

NODIA

processing units.

279. Vector processing is described as being most efficient for applications involving:
(A) extensive conditional branching and decision making.
(B) frequent input-output interactions with peripheral devices.
(C) operations on large, one-dimensional arrays of data.
(D) recursive subroutine calls and stack-based operations.

280. Which of the following best describes the functioning of an instruction pipeline?
(A) A single instruction is executed by multiple processors simultaneously.
(B) The processor executes instructions in a random, non-sequential order.
(C) Instructions are batched together and executed as one large block.
(D) Multiple instructions are in different stages of execution at the same time.

Direction : Please read the following source carefully.

The Input-Output (I/O) Organization of a computer system manages the communication between the central processor and external peripheral devices. This communication is mediated by an I/O interface, which resolves the differences in speed, data format, and control signals between the CPU and devices like keyboards, printers, and disks. A common challenge is timing, often addressed through asynchronous data transfer, which uses handshaking signals (e.g., strobe, ready, acknowledge) to coordinate data exchange without a common clock. There are several modes of transfer for moving data. Programmed I/O is the simplest, where the CPU is fully occupied checking the status of the peripheral and managing the transfer. A more efficient method is interrupt-driven I/O, where the peripheral signals the CPU with an interrupt when it is ready, freeing the CPU to perform other tasks. For high-volume data transfers, Direct Memory Access (DMA) is the most efficient mode. A DMA controller takes over the system bus to transfer blocks of data directly between a peripheral device and main memory, without continuous CPU intervention. When multiple devices can generate interrupts, a priority interrupt system is needed to determine which interrupt to service first. Serial communication is a method for transferring data one bit at a time over a single channel, common for long-distance communication.

Based on the above source, attempt the following five questions.

281. What is the primary function of an Input-Output Interface?
(A) To resolve the differences between the CPU and peripheral devices.
(B) To perform arithmetic calculations on behalf of the CPU.
(C) To store programs and data for long-term archival.
(D) To execute vector processing operations on arrays of data.

282. Which mode of transfer requires the most continuous CPU involvement for its entire duration?
(A) Direct Memory Access (DMA)
(B) Programmed I/O
(C) Interrupt-driven I/O
(D) Asynchronous I/O

283. In the context of I/O, what is the main advantage of Direct Memory Access (DMA) over interrupt-driven I/O?
(A) It simplifies the I/O interface hardware required for peripherals.
(B) It provides a highly efficient method for transferring large data blocks without CPU intervention.
(C) It uses serial communication to increase data transfer speed over short distances.
(D) It eliminates the need for a priority interrupt system.

284. Asynchronous data transfer relies on handshaking signals to:
(A) coordinate the exchange of data when a shared clock is not used.
(B) broadcast a common clock signal to all connected devices.
(C) select the highest priority device during a DMA transfer.
(D) check for errors in the data using a parity bit.

Click Here to Buy book on Amazon

285. A priority interrupt system is necessary when:
(A) a single device needs to transfer a large block of data via DMA.
(B) serial communication is used instead of parallel communication.
(C) multiple devices may request the CPU's attention simultaneously.
(D) the CPU is using programmed I/O to check a device's status.

Direction : Please read the following source carefully.

A computer's memory system is organized as a hierarchy to provide a balance between speed, capacity, and cost. At the top of the hierarchy is Cache Memory, which is a small, extremely fast memory that stores frequently accessed data and instructions from the Main Memory. It operates on the principle of locality of reference, bridging the significant speed gap between the fast CPU and the slower, larger Main Memory (typically DRAM). Below main memory is Auxiliary Memory (e.g., magnetic disks, solid-state drives), which is non-volatile and provides vast storage capacity at a lower cost. To further optimize memory usage, the concept of Virtual Memory is employed. This technique gives a program the illusion of having a very large, contiguous main memory space, when in reality its data and instructions may be fragmented and partially stored in auxiliary memory. The Memory Management Hardware (MMU), a component of the CPU or its chipset, handles the translation between the program's logical addresses and the physical addresses in main memory. A specialized, content-addressable memory known as Associative Memory is sometimes used in this process, particularly for rapidly searching page tables in virtual memory systems or for cache management, as it allows for searching based on data content rather than by address.

Based on the above source, attempt the following five questions.

286. What is the primary purpose of Cache Memory in the memory hierarchy?
(A) To provide long-term, non-volatile storage for user files.
(B) To create the illusion of a larger main memory than is physically present.
(C) To translate logical program addresses into physical memory addresses.
(D) To bridge the speed mismatch between the CPU and Main Memory.

287. The concept of Virtual Memory is designed to:
(A) increase the speed of auxiliary memory devices.
(B) eliminate the need for a memory hierarchy by using only one type of memory.
(C) allow programs to use a larger address space than the physical main memory.
(D) enable data to be searched by its content instead of its address.

NODIA

288. According to the passage, the operation of cache memory is effective due to which principle?
(A) The principle of asynchronous data transfer.
(B) The principle of locality of reference.
(C) The principle of last-in, first-out (LIFO) access.
(D) The principle of microprogrammed control.

289. What is the distinguishing feature of Associative Memory?
(A) It can only be accessed sequentially.
(B) It is the slowest but largest memory in the hierarchy.
(C) It is a volatile memory that requires constant power to retain data.
(D) It allows for data retrieval based on content rather than address.

290. The responsibility of translating a program's logical addresses to physical addresses in a virtual memory system falls to the:
(A) Memory Management Hardware.
(B) Auxiliary Memory.
(C) Cache Memory.
(D) Associative Memory.

Direction : Please read the following source carefully.

Multiprocessors are computer systems that contain two or more central processing units (CPUs) sharing access to a common main memory and peripherals. The primary characteristic of multiprocessors is the ability to execute multiple instruction streams in parallel, leading to significant increases in throughput and computational power. The effectiveness of such a system heavily depends on its Interconnection

Structure, which dictates how the processors are connected to each other and to the shared memory. Common structures include the time-shared common bus, the crossbar switch, and multistage switching networks, each offering different trade-offs in terms of cost, complexity, and performance. When multiple processors attempt to access a shared resource like the bus or memory simultaneously, a mechanism for Interprocessor Arbitration is required to resolve the conflict and grant access to one processor at a time. Furthermore, processors must coordinate their activities through Interprocessor Communication and Synchronization. A critical challenge in shared-memory multiprocessors is maintaining Cache Coherence. This problem arises because each processor typically has its own local cache. If one processor modifies a data item in its cache, other processors holding a copy of the same data in their caches will have stale, inconsistent information. Protocols are required to ensure that all caches and main memory present a consistent, coherent view of the data. The trend towards Multicore Processors integrates multiple processor cores onto a single chip, making these multiprocessor concepts relevant even in desktop and mobile computing.

Based on the above source, attempt the following five questions.

291. What is a defining characteristic of a multiprocessor system as described in the passage?
(A) Each processor has its own private, unshared main memory.
(B) It relies on a single, powerful processor with a large cache.
(C) It contains multiple CPUs that share access to a common memory.
(D) It is designed exclusively for serial, non-parallel task execution.

292. The Cache Coherence problem in a multiprocessor system refers to the issue of:
(A) maintaining data consistency across multiple private caches.
(B) ensuring that all processors run at the same clock speed.
(C) resolving simultaneous requests for a shared bus.
(D) connecting processors and memory through an interconnection structure.

NODIA

293. What is the function of an Interconnection Structure like a crossbar switch?
(A) To ensure cache coherence among all processors.
(B) To arbitrate access when multiple processors request a resource.
(C) To synchronize the execution of instructions across different cores.
(D) To facilitate communication paths between processors and memory.

294. Interprocessor Arbitration is a mechanism required to:
(A) translate logical addresses to physical addresses.
(B) resolve conflicts when multiple processors request a shared resource.
(C) manage the execution of a single task in parallel.
(D) maintain consistency of data in multiple caches.

295. The passage suggests that Multicore Processors are essentially:
(A) a type of processor that does not require cache memory.
(B) systems that are immune to the cache coherence problem.
(C) an implementation of multiprocessor architecture on a single chip.
(D) an alternative to interconnection structures.

Direction : Please read the following source carefully.

Two contrasting philosophies in CPU design are the Complex Instruction Set Computer (CISC) and the Reduced Instruction Set Computer (RISC). The CISC approach aims to close the semantic gap between high-level programming languages and machine language by providing a large and rich instruction set. CISC instructions can be very powerful, performing complex, multi-step operations (e.g., multiply and add to memory) in a single instruction. This design philosophy often involves a wide variety of addressing modes and variable-length instruction formats. To manage this complexity, CISC processors typically rely on a microprogrammed control unit, where control logic is stored in control memory. In contrast,

the RISC philosophy emphasizes simplicity and speed. A RISC computer features a small, highly optimized set of instructions that are generally of fixed length and execute in a single clock cycle. Complex operations are not performed by single instructions but are synthesized by the compiler from a sequence of simple ones. RISC architectures strictly follow a load/store model, meaning only explicit load and store instructions can access memory; all other operations work on data held in a large number of registers. This simplicity facilitates a faster, hardwired control unit and makes RISC processors exceptionally well-suited for pipelining, as the uniformity of instructions simplifies pipeline design and reduces stalls.

Based on the above source, attempt the following five questions.

296. Which of the following is a characteristic feature of a CISC architecture?
(A) A small, optimized set of single-cycle instructions.
(B) A hardwired control unit for maximum speed.
(C) An architecture where only load/store instructions access memory.
(D) Complex instructions that can perform multi-step operations.

297. The RISC design philosophy achieves complex operations by:
(A) having the compiler combine a sequence of simple instructions.
(B) using single, powerful instructions that access memory directly.
(C) relying on a microprogrammed control unit to sequence micro-operations.
(D) employing a large variety of complex addressing modes.

298. According to the passage, why are RISC processors particularly well-suited for pipelining?
(A) Because they use a microprogrammed control unit that can be easily modified.
(B) Because the uniformity and single-cycle nature of their instructions simplify pipeline flow.
(C) Because their complex instructions reduce the total number of pipeline stages.
(D) Because their instructions have variable lengths, allowing for dynamic pipeline adjustments.

299. A load/store architecture as described in the context of RISC means that:
(A) every instruction can access memory to fetch operands.
(B) all data must be loaded from memory into registers before being operated on.
(C) the processor has a very small number of general-purpose registers.
(D) instructions are loaded from a special, separate memory.

300. What type of control unit is typically associated with a CISC processor, and why?
(A) A microprogrammed unit, because it provides the flexibility to manage a complex instruction set.
(B) A hardwired unit, because it is faster for simple instructions.
(C) A pipelined unit, because CISC instructions are long and can be broken down.
(D) A load/store unit, because it is required for the numerous addressing modes.

Downloaded from www.pdf.tube

ANSWERS

1. (B) A logic gate
Fundamental components in digital electronics use physical devices to implement Boolean functions through electronic signals. These circuits act as basic building blocks by processing binary inputs to produce a single logical output. They form the core architecture for complex processors and memory systems, enabling the execution of computational tasks within modern hardware structures.

2. (A) Associative law
Mathematical principles in Boolean algebra define how logical operations interact within complex expressions. This specific law demonstrates that the grouping of variables within addition or multiplication does not change the final result. It allows designers to rearrange parentheses to simplify circuit diagrams while maintaining functional integrity, ensuring that sequential operations yield consistent binary outcomes.

3. (B) Both Statement I and Statement II are false.
Combinational logic relies exclusively on current inputs to determine output states, lacking internal storage mechanisms for past information. Conversely, sequential systems incorporate memory elements to track historical data alongside present inputs. Both statements incorrectly describe these categories by swapping their inherent characteristics regarding memory usage and the dependency of output generation within digital hardware systems.

4. (A) P-3, Q-2, R-1, S-4
Logic gates are mathematically represented by specific Boolean expressions that define their functional behavior. NAND operations produce the inverse of a product, while NOR functions invert a sum. XOR gates detect inequality between binary inputs, and XNOR circuits determine equivalence. These distinct mappings allow engineers to translate logical requirements into precise algebraic formulas during digital design processes.

5. (C) A larger group of 1s leads to a simpler term in the final expression.
Mapping techniques simplify Boolean functions by identifying adjacent cells containing specific logic states. Creating larger clusters of these states eliminates more variables from the resulting term, directly reducing the complexity of the hardware implementation. Efficient minimization relies on maximizing group sizes in powers of two to ensure the most concise mathematical representation of the original logic circuit.

6. (B) 2^n
Decoding units translate coded inputs into a singular active output among multiple possibilities. The relationship between input lines and potential outputs is exponential, as each additional bit doubles the number of unique combinations that can be represented. This architectural standard ensures that binary codes can uniquely address or activate specific hardware components or memory locations within a digital computer system.

7. (D) AND Flip-Flop
Sequential logic circuits utilize specific bistable devices to maintain state information over time. Standard types include variants that set or reset, toggle their state, or delay data until a clock pulse occurs. While various logic functions exist as simple gates, they do not all function as memory-storing elements, which are necessary for constructing registers and counters in processors.

8. (B) (a), (b) and (c) only
Sequential systems utilize internal storage to ensure that outputs reflect both current signals and historical data. Asynchronous variants operate based on input changes rather than relying on a global timing signal. While these memory-based components define sequential logic, registers serve as primary examples of this category, distinguishing them from combinational circuits which lack any capability for information storage.

NODIA

9. (C) 2^n
Data selection components utilize specific control signals to determine which binary input is routed to a single output. The capacity for selection is determined by the number of control lines, where each line adds a binary dimension to the possible choices. This exponential relationship allows a small number of selectors to manage a significantly larger array of incoming data.

10. (A) A group of flip-flops used to store a binary word.
Digital storage within a processor utilizes specialized circuits to hold temporary data for immediate processing. These units consist of multiple memory elements arranged in parallel to accommodate multi-bit binary strings. By capturing and maintaining bit patterns until explicitly updated, they provide high-speed access to operands, enabling efficient instruction execution and data manipulation within the central processing unit.

11. (D) Asynchronous counters are faster than synchronous counters for large numbers of bits.
Timing characteristics differ significantly between various counter architectures used in digital systems. Synchronous designs trigger all stages simultaneously, minimizing propagation delays and allowing for higher operational frequencies. In contrast, ripple designs experience cumulative delays as signals pass through each stage sequentially. Therefore, synchronous circuits offer superior performance as the number of bits in the counter increases.

12. (B) b, d, c, a
Designing logical hardware involves a systematic progression from abstract requirements to physical implementation. The process begins by defining functional behavior in a table, followed by deriving mathematical expressions. These expressions are then minimized to reduce component count before the final circuit diagram is constructed. Following this specific order ensures that the resulting hardware accurately and efficiently performs intended tasks.

13. (C) To store binary information in the form of bits.
Storage components serve as the repository for data and instructions within a computing environment. They maintain information in discrete states, allowing the system to retrieve or modify specific values as required by the processor. This capability is essential for preserving the results of computations and holding the software programs that dictate the computer's functional behavior over time.

14. (D) The development of integrated circuits (ICs)
Advancements in semiconductor manufacturing have enabled the placement of thousands of transistors onto single small chips. This technological evolution drastically reduced the physical size of computing hardware while simultaneously increasing logical complexity and processing power. By eliminating bulky discrete components, these circuits provided the foundation for modern portable devices and high-performance processors used in contemporary digital systems.

15. (A) Both A and R are true and R is the correct explanation of A.
Algebraic simplification rules allow for the reduction of Boolean terms without changing the functional output. The specific law mentioned describes how redundant variables are removed when combined with their complements in specific sum-of-product forms. Applying this principle streamlines logical expressions, which leads to more efficient circuit designs with fewer required gates, maintaining the same underlying mathematical logic.

16. (B) AB + BC + AC
Logic circuits designed for decision-making often require a specific threshold of active signals to produce an output. In this scenario, the expression captures all combinations where at least two out of three variables are high. This sum of products ensures that the system identifies a majority state, providing a robust method for error filtering or voting mechanisms in digital hardware.

17. (B) AND - NAND
Binary operations often have direct opposites that produce an inverted result for the same input combinations. One gate performs a logical conjunction, while its counterpart applies a negation to that result. Understanding these pairings is crucial for hardware design, as it allows

Click Here to Buy book on Amazon

engineers to use universal gates to implement any standard Boolean function through the application of De Morgan's theorems.

18. (A) Both Statement I and Statement II are true.
Functional components in digital logic serve specific roles in data routing and signal translation. One device interprets encoded binary patterns to trigger a unique output line, facilitating memory addressing or instruction interpretation. The other acts as a data selector, choosing between multiple input sources to feed a single destination. Both devices are fundamental to the architecture of modern processors.

19. (C) (a), (b), (c), and (d) only
Bistable storage elements form the basis of sequential logic by maintaining two stable states. Specific designs, like the JK type, may experience timing issues when inputs remain high for too long. These components can be converted between types using inverters and exhibit unique behaviors like toggling. Unlike combinational circuits, they possess memory, which is essential for storing state information.

20. (D) To miniaturize and combine a large number of electronic components into a single chip.
Microelectronics technology allows for the integration of numerous transistors, resistors, and capacitors onto a single semiconductor substrate. This consolidation reduces the physical footprint of electronic devices while improving processing speed and reliability. By housing entire functional circuits within a compact package, these components enable the sophisticated design of modern computing systems and small-scale mobile hardware devices.

21. (C) Solve problems by processing information in discrete form.
Modern computing machines operate by manipulating finite sequences of symbols, typically represented as binary digits. This discrete approach allows for precise calculations and logical operations that are distinct from continuous signal processing. By using fixed states to represent data, these machines can reliably store, retrieve, and transform information to execute complex algorithms and solve diverse computational problems efficiently.

22. (B) Counter
Logic circuits are categorized based on whether they possess internal memory. Devices like adders, decoders, and multiplexers generate outputs based solely on their current inputs, making them combinational. However, units designed to track sequences or count events must store their current state to determine the next one, classifying them as sequential circuits rather than part of the combinational group.

23. (A) P-4, Q-1, R-2, S-3
Specific digital components are defined by their primary roles within a computer's architecture. Arithmetic units calculate sums, while selectors route data from various sources. Translation units convert binary codes into distinct signals, and storage elements hold information for future use. Mapping these circuits to their functions provides a clear understanding of how hardware manages data flow and numerical computation.

24. (A) 1011
Binary counting devices progress through a sequence of states in response to external timing signals. Each pulse increments the stored value by one, following the standard rules of binary arithmetic. Starting from a zero state, a specific number of pulses will result in a binary representation equivalent to that decimal count, assuming the counter has sufficient capacity to hold the value.

25. (C) The octal number system uses digits from 0 to 8.
Positional number systems utilize a specific radix to define the range of symbols available for representing values. In a system based on eight, the valid digits range from zero to seven. Including the radix value as a single digit is mathematically incorrect, as that value is represented by a multi-digit sequence starting with one, following the standard rules of numerical representation.

26. (A) A representation where the position of the radix point is fixed.
Numerical data formats in computing often designate a specific, unchanging location for the separator between whole and fractional parts. This approach simplifies hardware design for arithmetic units, as the decimal or binary point does not shift during calculations. It is commonly

NODIA

used for integer representation or financial data where a constant precision level is required for consistent mathematical results.

27. (D) 22.25
Converting binary sequences to decimal involves summing the values of each bit based on its power-of-two position. Integer bits represent increasing positive powers of two, while fractional bits represent decreasing negative powers. Adding the calculated values for each active bit position yields a precise decimal equivalent, allowing for the translation of machine-readable binary strings into human-readable numerical values.

28. (A) Both Statement I and Statement II are true.
Signed number representation in digital systems utilizes a specific complementary form to simplify hardware arithmetic. By adding one to the bitwise inverse of a number, computers can treat subtraction as an addition operation. This method eliminates the need for separate subtraction logic, allowing a single adder circuit to handle both types of calculations efficiently while managing negative values.

29. (A) Both A and R are true and R is the correct explanation of A.
Scientific notation for binary data enables the representation of extremely large or small values by partitioning bits into separate components. This structure uses a significant part and a scaling factor, which allows the decimal point to move dynamically. Consequently, this format provides a much greater dynamic range than fixed-point systems, making it essential for complex engineering and physics simulations.

30. (D) The sign of the number
Binary data formats for signed integers reserve a specific bit position to indicate whether a value is positive or negative. Typically, the most significant bit serves this purpose, where a zero represents positive values and a one indicates negative values. This convention allows the arithmetic unit to correctly interpret and process numerical data during addition, subtraction, and comparison operations.

31. (A) It can detect the occurrence of an odd number of bit errors.
Data integrity during transmission is often maintained by adding a redundant bit to ensure a consistent count of set bits. This method allows the receiving system to identify if a single bit has flipped due to noise or interference. While it cannot correct errors or detect an even number of changes, it provides a simple, low-overhead diagnostic tool.

32. (D) (a), (b), (c) and (d)
Binary arithmetic follows specific rules for combining bits, mirroring decimal addition but with a base of two. When two ones are combined, the result is zero with a carry to the next position. If three ones are added, including a carry, the result is one with a new carry. These fundamental rules govern how all digital adders perform calculations.

33. (B) Floating-point representation
Complex calculations involving vast differences in scale require a data format that can adapt its precision. By separating a number into its significant digits and an exponent, computers can represent values ranging from subatomic measurements to astronomical distances. This flexibility makes it the standard choice for scientific software and high-precision engineering applications where fixed-point ranges are insufficient.

34. (A) A + (2's complement of B)
Digital processors simplify arithmetic by converting subtraction tasks into addition problems using a specific numerical transformation. By inverting the bits of the subtrahend and adding one, the computer creates a value that, when added to the minuend, yields the correct difference. This efficient approach allows the same hardware circuitry to perform both basic arithmetic operations seamlessly.

35. (C) The sign bit applies only to the exponent.
Scientific binary formats utilize a specific structure to represent real numbers, including bits for the fraction and the power. A dedicated bit determines whether the entire value is positive or negative, which is distinct from the sign of the scaling factor itself. Misunderstanding this part leads to incorrect data interpretation, as the overall value's sign is a critical attribute.

Downloaded from www.pdf.tube

36. (B) P-2, R-1, S-3, T-4
Error management in digital communication utilizes various techniques to ensure data reliability. Simple bits detect single changes, while more complex codes can identify and fix specific errors. Polynomial calculations and block-summing methods provide robust protection for large data transfers. Matching these methods to their characteristics is essential for understanding how computers maintain accurate information across noisy channels.

37. (C) Character strings
Arithmetic units in a processor are optimized to handle numerical data types that map directly to binary hardware structures. While integers, floating-point values, and coded decimal formats are processed through dedicated mathematical circuits, text-based data is typically managed through memory movement and comparison operations. These non-numeric types are represented by numbers but are not themselves operands for standard arithmetic functions.

38. (B) The sum requires more bits than the allocated storage.
Calculating results that exceed the capacity of a fixed-size register causes a specific error condition in digital arithmetic. This occurs when the carry out of the most significant position cannot be accommodated, leading to an incorrect stored value. Hardware must detect these events to ensure the accuracy of computations, especially when performing operations on signed or large unsigned integers.

39. (A) Both Statement I and Statement II are true.
Different numerical representations handle the value of zero in unique ways. Some systems produce two distinct bit patterns for positive and negative zero, which can complicate logical comparisons. Modern architectures prefer a unique representation to simplify zero-detection circuits and improve mathematical consistency. Both statements accurately reflect these architectural differences in how sign-magnitude or complementary systems manage the null value.

40. (B) (a), (b), (c), and (d) only
Numerical systems utilize positional notation where symbols represent values based on their location within a sequence. Hexadecimal uses sixteen symbols, while octal and binary follow specific radix-based conversions involving grouping or division. These properties are fundamental to computer science, though binary systems remain positional, unlike non-positional systems which do not change digit values based on their relative placement.

41. (C) They can detect the position of a single-bit error and correct it.
Advanced error-control techniques use multiple redundant bits to create a sophisticated map of a data block. Unlike simple methods that only signal an error, these codes identify exactly which bit flipped during transmission. This capability allows the system to restore the original data without requesting retransmission, significantly improving efficiency in storage devices and high-speed communication networks.

42. (A) -11
Signed integer interpretation requires converting a binary sequence using specific mathematical rules. For negative values starting with a leading one, the magnitude is found by inverting the bits and adding one before applying a negative sign. This conversion reveals the precise decimal value represented by the bit pattern within the constraints of an eight-bit complementary storage format.

NODIA

43. (D) The sequence of microoperations and the registers involved.
Hardware description methods utilize a symbolic notation to represent the internal data flow of a computer. This language focuses on how information moves between storage units and the specific elementary operations performed during each clock cycle. It provides a formal way to document the processor's architecture and the execution path of various machine instructions without physical circuit diagrams.

44. (A) Both A and R are true and R is the correct explanation of A.
Maintaining data integrity is vital because hardware components and transmission lines are susceptible to environmental interference and physical wear. These factors can unexpectedly alter binary values, leading to software errors or system failures. Implementing mathematical

checks allows systems to verify information accuracy, ensuring that the data processed or stored remains identical to its original intended state.

45. (C) Booth's algorithm - Signed integer multiplication
Specific mathematical procedures are implemented in hardware to optimize arithmetic tasks. One well-known method streamlines the process of multiplying binary numbers that include a sign bit, making it more efficient than traditional repeated addition. Identifying the correct application of these algorithms is key to understanding how processors achieve high performance when executing complex numerical instructions on signed data.

46. (D) Handle multiplication of signed binary numbers efficiently.
Conventional multiplication methods often struggle with negative values, requiring complex adjustments. A specialized algorithmic approach addresses this by treating strings of ones as a single operation, which reduces the total number of additions and subtractions needed. This efficiency is particularly valuable for hardware multipliers, as it speeds up processing for both positive and negative integers in modern computers.

47. (A) To convert a decimal integer to binary, the number is repeatedly multiplied by 2.
Changing numerical representations involves specific mathematical procedures based on the target base. For whole numbers, the process requires repeated division by the radix to extract remainders, whereas fractional parts are handled through multiplication. Confusing these methods leads to incorrect results, as the direction of the operation must match the part of the number being converted.

48. (B) Comparison, shifting, and subtraction
Digital division mirrors the manual long-division process by iteratively testing how many times a divisor fits into a dividend. The hardware performs this by aligning the numbers, checking their relative values, and removing the divisor when possible. This repetitive cycle of testing and adjusting bit positions is the standard way processors calculate quotients and remainders for integer arithmetic.

49. (B) An elementary operation performed on data stored in registers.
Processor activity is broken down into the smallest possible functional steps that occur within a single clock pulse. These actions include transferring data between internal storage units, incrementing counters, or performing basic bitwise logic. By combining these simple tasks, the control unit builds the complex sequences required to fetch, decode, and execute full machine-level instructions during operation.

50. (D) Statement I is false, but Statement II is true.
Data coding schemes serve different purposes in digital systems. One method represents decimal digits using a standard four-bit binary sequence, which is essential for certain arithmetic tasks. Another scheme ensures only one bit changes between consecutive values, though it does not assign specific weights to bit positions. Understanding these distinctions is crucial for correctly implementing input devices and numerical displays.

51. (C) The contents of R1 are transferred to R2.
Symbolic notation in computer architecture defines the direction of data movement between storage elements. The arrow indicates that the information residing in the source unit is copied into the destination unit, overwriting its previous value. This fundamental operation is the basis for all higher-level data processing tasks, as information must be moved to registers before execution.

52. (A) To provide a shared communication path between multiple registers and units.
Efficiency in hardware design is achieved by utilizing a centralized set of wires to connect various internal components. Instead of having dedicated connections between every possible pair of storage units, this shared infrastructure allows any unit to send data to another. This approach significantly simplifies the physical wiring of the processor while maintaining flexible and robust data transfer capabilities.

53. (A) Both Statement I and Statement II are true.

Click Here to Buy book on Amazon

Internal processor tasks are divided into two main categories based on the nature of the data involved. One group handles numerical calculations like addition and subtraction, while the other focuses on manipulating individual bits within a word for non-numeric purposes. Both types are essential for the comprehensive functionality of a computer, enabling both mathematical computation and logical data processing.

54. (C) P-2, Q-1, R-4, S-3
Symbolic representations in hardware description languages provide a concise way to document specific data movements and transformations. These include shifting bits to the left or right, inverting values, and resetting storage units to zero. Matching these symbols to their corresponding physical actions is fundamental for engineers designing and analyzing the micro-architecture of a central processing unit.

55. (C) It is denoted by DR ← M[AR].
Retrieving information from the main storage unit requires specifying the target location and a destination for the data. The notation indicates that the word located at a specific address is moved into a data-holding register. This operation is a critical component of the instruction cycle, as the processor must read data from memory before it can perform calculations.

56. (A) 0110
Bitwise operations compare corresponding positions in two binary strings to produce a new result. This specific logical function yields a high output only when the input bits are different. By applying this rule to each pair of bits in the given registers, the resulting pattern represents the combined state, illustrating how processors perform non-numeric data transformations internally.

57. (D) Arithmetic
Bit manipulation operations include several methods for moving data within a register. While some focus on simple logic or rotating bits back to the start, another specialized type is designed to preserve numerical properties. This third method handles the sign bit specifically during right shifts to ensure that the mathematical integrity of signed binary numbers is maintained during the operation.

58. (D) (a), (b), (c), and (d)
Shared communication paths in a processor are efficiently managed through specialized hardware like multiplexers and buffers. This design minimizes physical connections while allowing the control unit to select specific sources and destinations for data. Although only one transfer occurs at a time, the flexibility provided by this architecture is essential for organized and complex register-to-register information movement.

59. (B) An elementary operation performed during one clock pulse on data in registers.
Internal processing is structured as a series of tiny, discrete actions synchronized by the system clock. Each of these basic tasks represents the smallest unit of work the hardware can perform, such as moving a value or incrementing a counter. These individual steps are combined by the control unit to fulfill the requirements of more complex machine-level instructions.

60. (B) R1 + R2 (Represents the addition operation)
Formal hardware descriptions require specific syntax to distinguish between operations and the resulting data transfers. Simply listing an arithmetic sum without a destination is an incomplete representation of a microoperation. Correct notation must always specify the register where the result is stored, ensuring that the movement and transformation of information are clearly defined within the architecture.

61. (C) P: R2 ← R1
Control functions in hardware description languages are used to specify the conditions under which a data transfer occurs. The colon separates the Boolean condition from the actual operation, indicating that the transfer only happens when the signal is active. this structured approach allows designers to precisely map logic gates to the sequences of operations performed by the processor.

62. (C) A is true but R is false.
Utilizing a centralized communication path significantly reduces the total number of wires needed in a multi-register system. While this

Downloaded from www.pdf.tube

design simplifies the architecture, the reasoning provided incorrectly suggests that it increases complexity by using direct connections for every path. In reality, the bus structure avoids these individual connections, making the overall system more organized and easier to manufacture.

63. (B) Slower overall transfer speed due to shared access.
Adopting a shared communication infrastructure introduces a bottleneck because only one data transfer can occupy the path at any given time. While this significantly reduces the physical complexity of the processor's internal wiring, it prevents simultaneous information movement between multiple pairs of registers. This sequential limitation is the primary trade-off for the simplified hardware architecture of common bus systems.

64. (C) n full-adders.
Building a multi-bit circuit for mathematical operations requires repeating a basic functional unit for each bit position. A device that can sum two bits along with a carry from a previous stage provides the necessary logic for integer arithmetic. Connecting these units in a series allows the processor to handle binary numbers of any length, forming a complete adder.

65. (A) Both Statement I and Statement II are true.
Bitwise movement operations vary based on how they handle the vacated positions at the ends of a register. One standard method always inserts a null value when shifting. Another specialized method for signed numbers ensures the most significant bit is duplicated to keep the value's sign consistent. Both statements accurately describe these fundamental data manipulation techniques used in processors.

66. (B) (a), (c), and (d) only
Processor tasks are categorized by the type of functional units they utilize. Actions that perform bitwise comparisons like conjunctions, disjunctions, or inversions are classified as logic-based. While addition involves mathematical carry-overs and shifts involve repositioning, these are distinct from the pure bit-manipulation tasks. Identifying these categories is essential for understanding the internal organization of the arithmetic and logic unit.

67. (B) 1101
Specialized right-shift operations for signed numbers are designed to maintain the numerical value's polarity. When the bits move toward the less significant positions, the original leftmost bit is copied into the new vacancy. This ensures that a negative value remains negative, allowing the processor to perform division by two while correctly preserving the sign of the binary integer.

68. (D) Selective-set - Sets bits in register A where corresponding bits in B are 1.
Bitwise manipulation instructions allow for precise control over the data stored in registers. One common operation uses a mask to force specific bits to a high state while leaving others unchanged. This is achieved by performing a logical OR with another register. Correctly identifying these operations is vital for software and hardware tasks that require bit-level data adjustment.

69. (C) A control function or signal
Hardware description syntax uses specific characters to represent the logic that triggers a register transfer. The variable placed before the colon acts as a gatekeeper, representing a physical signal from the control unit. The operation only proceeds when this signal evaluates to a high state, ensuring that data movement is synchronized with the overall system timing and logic.

70. (B) To act as a switch that can disconnect a device from the bus.
Managing shared communication lines requires components that can effectively isolate inactive units. This specialized device can enter a high-impedance state, effectively removing its electrical connection from the circuit. This allows multiple sources to be connected to the same wires without interference, as only one unit is permitted to drive the signal at any given time.

71. (C) Instructions and data are stored together in the same memory.
This foundational architectural concept allows a processor to treat both computational tasks and information as binary sequences in a unified storage space. By doing so, the machine can easily modify its own operations and access different programs without physical reconfiguration. This flexibility is what enables modern computers to

NODIA

function as general-purpose devices capable of executing diverse software applications.

72. (C) Program Counter (PC)
Sequencing instructions requires a dedicated storage unit to track the current location in the program code. This register automatically updates after each fetch to point to the subsequent command, ensuring the processor follows the intended logical flow. It is a vital component of the control unit, enabling the computer to execute complex code structures including loops and conditional branches.

73. (C) Statement I is true, but Statement II is false.
Instruction formats consist of multiple parts that define the machine's next action. One section specifies the specific task to be performed, such as addition or data movement. However, other sections are used to identify where the data is located, rather than the task itself. Both statements accurately distinguish between the instruction's purpose and its internal functional components.

74. (A) b, a, d, c
Executing a machine command follows a rigid chronological sequence of events. The processor first retrieves the code from memory and interprets its meaning. If the instruction references a memory location indirectly, it must then perform an additional read to find the final address. Only after these preparatory steps are complete can the actual functional task be performed by hardware.

75. (C) To generate timing signals for sequencing all operations.
Orchestrating the internal operations of a processor requires precise synchronization across various hardware components. This specialized unit produces a continuous stream of pulses that dictate the exact moment when data transfers or microoperations occur. By managing the temporal flow of events, it ensures that instruction execution phases proceed in the correct order, maintaining the overall stability of computation.

76. (C) The address of the memory word that contains the address of the operand.
Complex memory access methods allow for greater flexibility in software design by using pointers. In this mode, the instruction does not point to the data directly but instead specifies a location that holds the data's final destination. This multi-step process is essential for implementing advanced data structures and dynamic memory allocation within a computer's operating environment.

77. (D) Floating-point instructions
Basic computer architectures are often designed with a minimal set of instruction categories to simplify the control logic. These typically include operations that access memory, manipulate internal registers, or handle communication with external devices. While advanced processors include specialized commands for complex decimals, these are not usually part of the fundamental core set found in a simplified basic model.

78. (D) (a), (b), (c) and (d)
Managing unexpected events requires a hardware-driven process that pauses current tasks. The system automatically preserves the return location and redirects the processor to a specialized service routine. Once the external or internal signal is handled, the saved address allows the computer to resume its previous activity. This entire cycle is fundamental for responsive and stable computer operations.

79. (B) To manage the information flow between the computer and external devices.
Modern computers must interact with a variety of peripherals and respond to real-time events. A dedicated subsystem handles these interactions, ensuring that data moves correctly between the central processor and external hardware. By utilizing signals to pause operations when needed, this system allows the computer to multitask and provide efficient communication with users and other machines.

80. (C) Process one instruction.
The operational lifecycle of a processor is built upon a repetitive sequence of events for every command it receives. This cycle includes retrieving the command, understanding its requirements, and performing the specified action. By completing these steps millions of

Click Here to Buy book on Amazon

times per second, the computer executes entire programs and manages all the logical tasks required by the software it runs.

81. (D) All instructions must have an address part.
Instruction formats are designed to be flexible, as not every operation requires an external memory reference. Some commands perform tasks exclusively on internal registers or clear specific flags, eliminating the need for an address field. This efficiency allows the architecture to use different bit arrangements for various tasks, optimizing the use of memory and processing time during execution.

82. (A) P-4, Q-3, R-2, S-1
Specific internal storage units within a processor are assigned unique roles to facilitate the instruction cycle. Some hold the locations of data, while others track the next command or store the actual binary code being processed. Understanding the primary function of each register is essential for analyzing how a computer fetches, decodes, and executes information during its normal operation.

83. (B) A hardware signal from a peripheral device or an internal event.
Unplanned interruptions in program execution are triggered by external hardware or specific internal conditions that require immediate attention. These electrical signals force the control unit to save its current state and jump to a different set of instructions. This mechanism is vital for handling user input, managing hardware errors, and ensuring that the system remains responsive to real-world demands.

84. (B) Instruction Register (IR)
Once a command is retrieved from memory, it must be stored in a dedicated location where the control unit can analyze it. This specialized register holds the binary pattern representing the opcode and operands for the duration of the current cycle. This allows the decoder to generate the specific control signals needed to complete the task defined by that command.

85. (D) Statement I is false, but Statement II is true.
Memory access techniques vary in how they specify the location of data. One method uses the instruction's address field as the final destination, while another uses that field to find a second location where the data's address is stored. Both statements correctly highlight these distinctions, which are fundamental to how software interacts with a computer's physical and logical memory.

86. (A) Both A and R are true and R is the correct explanation of A.
Storing commands in a unified memory space with data revolutionized computing by allowing programs to be treated as information. This versatility enables the creation of sophisticated software features like branching and loops, as the processor can read and modify the instructions it executes. This principle serves as the core foundation for all modern programmable computer architectures and operating systems.

NODIA

87. (B) It fetches the data from the memory location specified by X.
In the simplest form of memory access, the address provided in the command points directly to the operand's location. The processor reads the value from that specific spot in the main storage and brings it into the accumulator or another register. This direct method is efficient for straightforward data retrieval where the location of the information is known beforehand.

88. (C) INP (Input Character) - Input-Output
Machine instructions are classified by the resources they interact with, such as memory, internal registers, or external devices. A command that retrieves data from a peripheral unit falls into a specific category designed for communication. Correctly identifying these pairs is essential for understanding the instruction set architecture and how the processor coordinates with the rest of the computer system.

89. (B) The physical address of the operand.
Determining where data actually resides often involves calculations or multiple memory reads, depending on the addressing mode used. The result of these steps is a final, absolute location that the processor uses to perform the read or write operation. This final destination is a critical concept in computer architecture, as it represents the true target of any memory-based instruction.

90. (C) (a), (b), (c), and (d) only

Traditional computer systems require a core set of hardware modules to function as general-purpose machines. This includes storage for programs, a unit to perform calculations, a control center for orchestration, and fast temporary storage. While networking is common today, it is not a fundamental requirement for the basic logic of a stored-program machine to execute and process binary instructions.

91. (A) AR ← PC, IR ← M[AR], PC ← PC + 1
Retrieving a command from memory involves a precise sequence of microoperations. The processor first identifies the location of the next command, then moves the data from that location into a temporary holding area. Finally, it increments its internal counter to prepare for the next step. This repetitive process is the first phase of every instruction the computer executes.

92. (D) They operate only on data present in CPU registers.
Commands that reference memory are specifically designed to interact with the larger storage units outside of the immediate processor registers. While they may use registers for results or addresses, their primary purpose is to move data to or from the main memory. Understanding this distinction is vital, as register-only commands are a separate category with different hardware requirements and speed characteristics.

93. (B) To enable the CPU to return to the correct point in the interrupted program.
When a system pause occurs, preserving the current state is essential for maintaining program continuity. By saving the location of the next intended instruction, the processor can resume exactly where it left off after the urgent task is finished. This mechanism ensures that the main software execution is not permanently disrupted by the hardware's need to respond to external signals.

94. (A) A binary code specifying a sequence of microoperations for the computer.
A machine instruction is essentially a coded set of signals that tells the hardware which elementary tasks to perform. These codes are interpreted by the control unit to trigger specific movements of data and logical transformations during each clock cycle. This low-level representation is the only language the processor truly understands and executes to perform complex computational work.

95. (C) To synchronize the timing of data transfer between the CPU and the I/O device.
Managing communication with external hardware requires signals to indicate when data is ready or when a task is complete. These indicators prevent the processor from reading data before it arrives or sending new information before the previous piece is processed. This synchronization is critical for maintaining accurate data flow between high-speed internal components and generally slower external peripheral devices.

96. (A) A low-level symbolic language that has a one-to-one correspondence with machine language.
This type of programming uses human-readable labels instead of raw binary strings to represent hardware instructions. Because each label maps directly to a specific processor command, it provides precise control over the machine while being easier for developers to write and debug. It serves as a vital bridge between high-level logic and the actual electronic operations of the computer.

97. (B) Translate a program from assembly language to machine language.
Software development for low-level tasks requires a tool to convert symbolic code into the binary format the hardware understands. This utility interprets mnemonics and labels, mapping them to their corresponding bit patterns and memory addresses. This conversion is a necessary step in the programming process, as it produces the final executable code that the central processing unit can run.

98. (A) Both Statement I and Statement II are true.
Computers operate exclusively on binary patterns, making raw machine code the only directly usable format for execution. To aid human developers, symbolic representations were created to replace complex bit strings with recognizable abbreviations. Both statements correctly describe the relationship between these two levels of coding, highlighting the necessity of binary execution and the readability benefits of symbolic mnemonics.

Downloaded from www.pdf.tube

99. (B) (b) and (c) only
The initial phase of converting symbolic code into machine format focuses on analyzing the structure and defining locations. This includes identifying all labels and their corresponding memory addresses to build a reference guide. It also involves interpreting directives that set the program's starting point or ending. This preparatory work is essential before the final binary instructions can be generated in subsequent steps.

100. (C) Sequence of instructions executed repeatedly until a condition is met.
Programming structures utilize repetitive execution paths to perform tasks on multiple data sets efficiently. These constructs continue to process a specific block of machine code until a predefined logical test evaluates to false. By automating repetitive operations, they reduce the need for redundant code entries and allow software to handle variable amounts of data through iterative processing cycles.

101. (A) An instruction that stores the return address.
Executing modular code blocks requires a mechanism to remember where the main program was paused. The system must record the location of the next instruction before jumping to the subordinate task. This saved information ensures that once the secondary task is complete, the processor can successfully find its way back to continue the original sequence of operations without error.

102. (C) ORG
Low-level code includes commands that the processor executes directly, such as adding numbers or moving data. However, it also contains instructions for the conversion tool itself that do not produce machine code. These directives guide the organization of the program in memory. Distinguishing between executable commands and these organizational instructions is fundamental for understanding how programs are built and loaded.

103. (C) P-4, Q-1, R-2, S-3
Directives in symbolic programming provide essential metadata for organizing code and declaring data types. They specify where in memory a program should start, where it ends, and the format of numerical constants used. Mapping these specific terms to their organizational roles is a key skill for managing how binary information is laid out and interpreted during the translation and loading process.

104. (C) To avoid code repetition and structure the program into modular blocks.
Efficiency in software design is achieved by isolating frequently used tasks into independent segments. This approach allows a single copy of code to be reused multiple times throughout a program, which conserves memory and simplifies debugging. By breaking complex logic into manageable parts, developers can create more organized, readable, and maintainable systems that perform complex operations reliably.

105. (B) A table that maps symbolic labels to their memory addresses.
During the conversion of symbolic code, the translation tool maintains an internal record of all names used to identify specific code or data locations. This resource allows the tool to replace human-friendly labels with the absolute numeric addresses required by the hardware. This mapping is essential for resolving branches and memory references, ensuring that the final binary code points to the correct locations.

106. (C) It does not require any special instructions; it uses standard data movement instructions.
Communicating with external devices often requires unique hardware commands that are distinct from standard memory or arithmetic tasks. These specialized instructions check for hardware readiness and trigger data transfers between the processor and peripherals. Understanding that standard data movement is often insufficient for these tasks is vital for correctly designing the input-output architecture of a computer system and its software.

107. (B) The CPU can only understand and execute instructions in binary machine code format.
Computers are electronic machines that operate based on the presence or absence of electrical signals, represented as zeros and ones. Because of this physical reality, any human-readable code must be transformed into these discrete patterns before the hardware can process it. This translation is the most fundamental requirement

NODIA

in computing, enabling the interaction between high-level logical concepts and physical circuitry.

108. (A) Both A and R are true and R is the correct explanation of A.
Writing code that maps directly to the machine's hardware allows for extreme optimization of memory and processing speed. By manually managing registers and instruction sequences, developers can eliminate the overhead often introduced by high-level languages. This fine-grained control is why low-level coding remains essential for performance-critical systems, such as embedded controllers and high-speed drivers, where resource efficiency is paramount.

109. (C) Translating instructions into machine code and resolving symbolic addresses using the symbol table.
The final step in creating an executable program involves converting all symbolic mnemonics into their final binary bit patterns. Using the reference guide created earlier, the tool fills in the actual memory locations for all labels. This comprehensive process produces the final machine language file that the processor can load and execute to perform the intended computational tasks defined by the programmer.

110. (A) Both Statement I and Statement II are true.
Unexpected events and structured task shifts both change the flow of execution in a processor. One is initiated by hardware signals from external sources, while the other is explicitly triggered by a program command. Both statements correctly identify the functional similarities in how they transfer control and the key difference in their underlying triggers, which is fundamental to computer architecture.

111. (D) Setting a register to -10 and incrementing it until it becomes zero.
Repetitive tasks are often managed using an internal counter that tracks completion. A common technique involves initializing a register with the negative value of the total desired iterations. After each cycle, the counter is increased, and the program checks for a zero state to determine if it should repeat or exit. This method is efficient for managing loops in low-level programming.

112. (C) Assembly Language - Low-level language
Programming environments are categorized by how closely they correspond to the underlying hardware architecture. Methods that use symbolic labels to represent specific machine instructions are classified as low-level because they offer direct control over the processor's registers and memory. This is distinct from high-level languages that provide abstract concepts and require extensive translation before they can be executed by the hardware.

113. (C) Binary or hexadecimal code
At the most fundamental level, computer commands exist as sequences of bits that trigger specific electronic actions. While often represented in base-sixteen for human convenience, these codes correspond directly to the electrical states within the processor. This raw format is the only one that the hardware can interpret and execute without further translation, representing the core logic of the machine.

114. (C) A dedicated hardware stack.
Managing the return path for nested code segments requires a robust storage mechanism that follows a specific order. A specialized structure allows the system to push return addresses as it enters new routines and pop them as it finishes. This hardware-supported approach ensures that the processor can always navigate back through multiple levels of execution, maintaining the stability of the program's logical flow.

115. (C) Programmed I/O is the most efficient method for high-speed data transfers.
Different data transfer methods offer varying levels of efficiency based on how they involve the central processor. Techniques that force the CPU to wait and repeatedly check device status are often slower and wasteful of resources. For high-speed transfers, other methods that bypass the main processor or use interrupts are far more effective, allowing the system to maintain overall performance.

116. (C) The microprogram consisting of microinstructions.
Control units in modern processors utilize a dedicated internal storage space for the low-level commands that implement machine instructions.

Click Here to Buy book on Amazon

This area holds the specific sequences of signals required to orchestrate the hardware's internal data flow. By storing these commands here, the architecture becomes more flexible, as the behavior of the entire processor can be modified by updating this internal firmware.

117. (A) Incrementing the control address register.
Progressing through the steps of a low-level routine requires a mechanism to identify the next command in the sequence. A dedicated storage unit tracks the current step, and increasing its value is the most common way to move to the subsequent microoperation. This fundamental capability allows the control unit to execute complex, multi-step instructions in a predictable and organized manner.

118. (A) Both Statement I and Statement II are true.
Processors can be controlled using either fixed physical circuits or stored sequences of low-level commands. While physical logic is faster due to its direct implementation, the use of stored programs offers greater design flexibility at the cost of some execution speed. Both statements accurately reflect these fundamental engineering trade-offs between speed and adaptability in different control unit architectures.

119. (D) (a), (b), (c) and (e)
Low-level command formats often include multiple fields to manage both hardware signals and the flow of control. These fields specify which internal operations to trigger, the conditions for branching, and where to find the next step in the sequence. By organizing bits in this way, the control unit can perform many tasks simultaneously, optimizing the execution of complex machine instructions.

120. (B) It involves writing a microprogram (firmware) that implements the computer's instruction set.
Designing an adaptable control unit relies on creating a specialized layer of software that translates high-level machine commands into hardware signals. This firmware acts as a bridge, defining how each instruction is executed through a sequence of smaller, elementary steps. This approach allows engineers to modify or expand the computer's capabilities without physically altering the underlying electronic circuits.

121. (B) The micro-cycle.
The internal operations of a control unit are structured as a repetitive process of retrieving and executing low-level commands. This cycle represents the smallest unit of time in which the control hardware performs its tasks, typically corresponding to a single system clock pulse. Understanding this foundational cycle is essential for analyzing the performance and timing characteristics of microprogrammed processor architectures.

122. (D) Main memory address for data operand.
Command formats within the control unit focus on orchestrating internal hardware signals and managing the sequence of microoperations. While they include fields for branching and triggering operations, they do not typically store the location of data in the computer's larger main memory. That information is instead managed by the instruction register and the address register during the overall instruction cycle.

123. (C) P-4, Q-3, R-1, S-2
Specific components within a microprogrammed control unit are assigned unique roles to manage the flow of low-level instructions. These include storage for the microprogram, registers to track the current step, and units that determine the next address in the sequence. Matching these components to their functional descriptions is essential for understanding how the control unit manages the execution of complex tasks.

124. (C) Greater flexibility to modify or add new instructions.
Using stored sequences of commands instead of fixed circuits allows designers to easily update the processor's behavior. This flexibility is the primary benefit, as new instructions can be implemented or existing ones optimized by simply changing the firmware. This approach significantly simplifies the design process and allows the hardware to adapt to changing software requirements or bug fixes over time.

125. (B) A sequence of microinstructions that implements a single machine instruction.
At the core of a processor's operation, each high-level command is broken down into a series of much smaller, elementary tasks. This collection of steps serves as the definitive guide for how the

NODIA

hardware should perform that specific action. By organizing these steps into structured routines, the control unit can consistently and accurately execute the entire set of instructions supported by the machine.

126. (C) The next address is always fetched from the main memory.
Determining the subsequent step in a low-level routine involves several possible methods, such as simple incrementing or branching based on internal conditions. Crucially, these addresses are retrieved from the specialized control storage within the unit, rather than the computer's large main memory. This internal retrieval is necessary for maintaining the high speeds required for processor control and synchronization.

127. (B) The number of microinstructions and the width of each microinstruction.
The capacity of the specialized internal storage for control commands is determined by both the total number of steps required and the amount of information contained in each step. A larger set of instructions or more complex parallel operations will necessitate a bigger storage unit. This design choice directly impacts the complexity and flexibility of the processor's control architecture and its capabilities.

128. (A) Both A and R are true and R is the correct explanation of A.
Microprogrammed architectures provide a significant advantage in adaptability because they decouple the instruction set from the physical hardware. By storing control logic as a program, engineers can modify the computer's functional behavior by simply updating the internal firmware. Both statements correctly highlight this relationship, explaining why this design choice is favored for its flexibility in building complex and evolvable processor systems.

129. (C) Control Address Register (CAR)
Within the control unit, a specific storage element is used to keep track of the current step in a microprogram. This register provides the address for the next low-level command to be retrieved from the control memory. By managing this address, the unit can correctly sequence the microoperations needed to execute any machine-level instruction, ensuring precise and reliable processor operation.

130. (A) Both Statement I and Statement II are true.
Microprogramming techniques offer different trade-offs between hardware complexity and operational parallelism. One approach uses many bits to trigger multiple tasks at once, while another uses fewer bits and a decoder to save space. Both statements accurately describe these two methodologies, highlighting the design choices engineers face when balancing efficiency and speed in the construction of a control unit's internal architecture.

131. (B) Any of the registers can be used as a source or destination for ALU operations.
Processor architectures that utilize a flexible set of internal storage units allow for more efficient data manipulation. In this design, instructions can specify any available register to hold operands or results, rather than relying on a single dedicated unit. This versatility simplifies the task of software compilers and reduces the need for frequent memory access, significantly improving overall system performance and speed.

132. (B) POP
Managing data in a specialized first-in, last-out structure requires specific commands for adding and removing items. To retrieve the most recently stored piece of information, a specific instruction is used that also updates the internal pointer to the next available location. This mechanism is fundamental for managing temporary variables, function calls, and complex mathematical expressions within a computer's processing environment.

NODIA

133. (A) Both Statement I and Statement II are true.
Processor designs have evolved into two primary philosophies regarding instruction complexity. One emphasizes simple, fast tasks that can be finished in a single step, while the other supports powerful, multi-step commands that can perform more work per instruction. Both statements correctly identify these architectural differences, which have significant impacts on hardware design, compiler technology, and the overall efficiency of program execution.

134. (A) b, a, c, d

Executing a machine command follows a rigid chronological sequence of events. The processor first retrieves the code from memory and interprets its meaning. If the instruction references a memory location indirectly, it must then perform an additional read to find the final address. Only after these preparatory steps are complete can the actual functional task be performed by hardware.

135. (A) The number of address fields in an instruction format depends on the internal organization of the CPU.
The way machine instructions are structured is directly tied to the underlying physical layout of the processor's registers and functional units. Systems that rely on a single accumulator will use different formats than those with many general-purpose registers or those that use a stack. This relationship ensures that the instruction set architecture is optimized for the hardware's specific capabilities.

136. (A) Indexed Addressing
Calculating the final location of data often involves combining a static value from the instruction with a dynamic value stored in a register. This method allows a single instruction to access different elements in an array by simply changing the register's content. It is a powerful technique for efficiently managing data structures and performing repetitive tasks on collections of information in memory.

137. (D) Immediate Addressing
Some instructions provide the required data value directly within the command itself, eliminating the need for any additional memory or register reads. This approach is the fastest way to load constants or perform simple arithmetic with fixed numbers. By including the operand in the bit pattern of the instruction, the processor can execute the task immediately after the decode phase is finished.

138. (C) (a), (b) and (d) only
Efficient processor designs often prioritize a large number of internal storage units and a simplified, uniform command set. By using hardwired control and avoiding complex, variable-length formats, these machines can execute tasks much more quickly and predictably. These characteristics are central to an architectural philosophy that focuses on high throughput and streamlined hardware, distinguishing it from more complex instruction set designs.

139. (B) To provide flexibility for accessing operands and managing data structures.
Different methods for identifying data locations allow programmers to write more efficient and compact code. These techniques enable direct access, pointers, and array indexing, which are essential for handling complex information. By providing various ways to calculate addresses, the architecture can support sophisticated software designs while maintaining high performance and making the most of the available memory and processor resources.

140. (B) The address of the top of the stack.
Managing a specialized storage area requires a dedicated register to track the current point of entry and exit. This pointer always identifies the location of the most recently added piece of data, allowing the system to easily add new items or retrieve existing ones. This mechanism is critical for the proper functioning of subroutines, interrupt handling, and the evaluation of complex logical expressions.

141. (C) They predominantly use a hardwired control unit for speed.
Architectures designed with complex and varied instruction sets typically rely on microprogrammed control rather than fixed physical circuits. This is because the wide range of instructions and addressing modes is too complicated to implement efficiently with simple logic gates. Microprogramming provides the necessary flexibility to manage this complexity, even though it may result in some reduction in the absolute speed of execution.

142. (C) P-2, Q-1, R-3, S-4
Different methods for determining where data is located involve various calculations and register references. These include using values directly from the instruction, referencing memory addresses, or combining values with the current program counter. Matching these modes to their mathematical definitions is essential for understanding how the processor retrieves

Downloaded from www.pdf.tube

information and manages the execution flow of a program across different memory locations.

143. (B) Cause: To minimize memory access operations. Effect: Faster program execution.
Increasing the number of internal storage units allows the processor to keep more frequently used data close to the functional units. This significantly reduces the need to wait for slower memory access, which is a major bottleneck in computing. By performing most tasks using these fast registers, the machine can achieve much higher overall performance and finish complex calculations in less time.

144. (C) Zero-address instruction format
Some machine commands do not require any explicit location for data because they operate on a specialized, implicit structure like a stack. In these cases, the instruction itself specifies the task, and the hardware automatically retrieves and stores information at the top of the stack. This results in very compact code, as no extra bits are needed for memory addresses or registers.

145. (A) Both Statement I and Statement II are true.
Advanced memory access techniques allow for more sophisticated data management in software. One method uses a register to point to a memory location, while another adds an offset to a fixed base value to find the final address. Both statements correctly describe these fundamental methods, which are essential for implementing dynamic data structures and efficiently navigating different regions of memory.

146. (A) Both A and R are true and R is the correct explanation of A.
Specialized processor organizations that use a stack-based structure are highly efficient at processing mathematical formulas arranged in a specific order. This notation allows the computer to perform calculations by simply moving operands and results on and off the top of the stack. Both statements correctly highlight this synergy, explaining why this design choice simplifies the hardware and logic required for complex arithmetic.

147. (A) A B + C D + *
Representing mathematical formulas without parentheses requires placing operators after their corresponding operands. This approach allows a computer to evaluate the entire expression in a single pass using a stack. By following the correct sequence, the processor can solve the inner additions before performing the final multiplication, illustrating how this logical notation is used to structure data for efficient hardware-based computation.

148. (C) PUSH A
Instruction formats describe how many operands an instruction uses. Three-address and two-address examples shown in options A and B are incorrect in structure, while the zero-address example is also mismatched. PUSH A correctly represents a single-operand stack instruction, where data from location A is pushed onto the stack, making option (C) the valid pair.

149. (B) Hold the address of the next instruction to be fetched.
Sequencing instructions requires a dedicated storage unit to track the current location in the program code. This register automatically updates after each fetch to point to the subsequent command, ensuring the processor follows the intended logical flow. It is a vital component of the control unit, enabling the computer to execute complex code structures including loops and conditional branches effectively.

150. (A) (a), (b), and (c) only
Standard instruction formats organize bits into specific functional areas to guide processor activity. The operation code defines the task, while the mode field determines how to interpret operand locations. Address fields specify where data resides in memory or registers. These three elements are essential for defining machine-level tasks, whereas metadata like authorship is irrelevant to the hardware execution logic of any computer.

151. (D) AC <- AC + M[X]
In machines that utilize a single primary working register, instructions often assume that one operand is already present in that register. A command that adds a value from memory will automatically use the contents of the accumulator and store the final result back into it. This approach simplifies the instruction

NODIA

format by reducing the number of bits needed to specify data locations.

152. (D) It guarantees that all programs will run faster than on an accumulator-based machine.
Utilizing a broad set of internal registers offers many advantages, such as reduced memory access and better organizational options for software. However, the final speed of any program depends on many factors beyond just the register layout, including clock frequency and the efficiency of the code itself. No single architectural feature can ensure superior performance across all possible scenarios or software designs.

153. (C) Decrementing the stack pointer, then writing data.
Managing a storage structure that expands toward lower numerical addresses requires a specific sequence of actions during data entry. The tracking register must first be decreased to point to the next available lower location before the new information is stored. This standard convention is common in many computer architectures and is essential for the organized management of the system's internal data stack.

154. (C) The complexity of the instruction set and control unit.
Computer architectures are primarily distinguished by their philosophical approach to command design and hardware control. One style uses a vast array of powerful, varied instructions managed by complex logic, while the other favors a small set of simple, fast commands with direct, hardwired control. This fundamental difference in complexity dictates how the processor is built and how software interacts with the machine.

155. (A) Both Statement I and Statement II are true.
Different ways to specify data in instructions offer varying levels of efficiency and flexibility. Some commands have the data built into their definition, while others include the actual value directly in the instruction code. Both statements accurately reflect these fundamental techniques, which are used to optimize program speed and reduce the need for external memory references during the execution of tasks.

156. (B) A method for specifying the location of the operands used by an instruction.
Determining where data is located is a critical part of executing any machine command. These various techniques provide the rules for calculating the final address of the required information, whether it is in a register, memory, or the instruction itself. By offering different methods, the computer can efficiently handle various programming tasks, such as accessing constants, variables, pointers, or large arrays.

157. (A) It would add the top two elements of the stack, pop them, and push the result.
In architectures that use a specialized first-in, last-out storage structure, instructions often operate on the data already at the top. An addition command automatically retrieves the two most recently added values, calculates their sum, and stores that result back onto the structure. This implicit approach eliminates the need for address fields, making the resulting machine code much more compact.

158. (C) Data manipulation instructions include arithmetic, logical, and shift operations.
Machine commands are grouped based on the specific type of task they perform on information. The largest category focuses on transforming data through mathematical calculations, bitwise logic, and repositioning bits. These tasks are distinct from simply moving information or controlling the flow of the program. Understanding these functional groups is essential for analyzing the capabilities and structure of any computer's instruction set.

159. (B) To support a wide variety of complex addressing modes and operations.
Some architectures provide a vast and flexible set of commands that can perform many different tasks and access data in numerous ways. To accommodate this versatility, instructions must vary in size, as simple tasks require fewer bits than those that involve multiple memory references or complicated calculations. This flexibility allows for powerful programming but requires more sophisticated hardware to decode and execute.

160. (B) Central Processing Unit
The primary hardware modules that manage the execution of programs and the processing of data are grouped together into a single core entity. This includes the logic for calculations,

Click Here to Buy book on Amazon

the management of internal storage, and the coordination of all tasks. Collectively, these components form the functional heart of the computer, responsible for interpreting instructions and directing the flow of information throughout the system.

161. (B) Using multiple processing units to solve a problem concurrently.
Increasing computational speed can be achieved by breaking tasks into smaller parts and executing them at the same time across different hardware. This approach moves beyond traditional sequential processing, allowing the system to handle vast amounts of data more efficiently. By coordinating multiple units to work on a single problem, modern computers achieve high performance and solve complex tasks much faster.

162. (C) Pipelining
Efficiency in processing is achieved by dividing a complex task into several distinct stages, allowing different instructions to be in various phases of completion at the same time. This works much like an assembly line, where multiple items are being worked on simultaneously by different stations. This technique significantly increases the total number of tasks finished per second, improving the throughput of the processor.

163. (A) Both Statement I and Statement II are true.
Different types of concurrent processing focus on improving specific aspects of a computer's operation. One method accelerates complex mathematical tasks by breaking down numerical calculations, while another overlaps the various phases of the instruction cycle. Both statements accurately describe these two applications, which are fundamental to increasing the overall speed and efficiency of modern high-performance central processing units.

164. (B) (a), (b) and (c) only
Several factors can disrupt the smooth flow of instructions through a multi-stage processor. These include conflicts over hardware resources, dependencies on data that has not yet been calculated, and the challenges of managing jumps and branches in the program flow. Understanding these hazards is essential for designing efficient systems that can maintain high performance despite the logical complexities of program execution.

165. (B) Performing the same operation on multiple data elements simultaneously.
High-performance computing often requires applying a single mathematical task to large collections of information, such as arrays or lists. By using specialized hardware that can process many values at once, the system achieves much higher speeds than by processing each item individually. This technique is particularly valuable for scientific simulations, engineering analysis, and advanced graphics processing tasks in modern computers.

166. (B) 12ns
In a multi-stage processing system, the overall timing is dictated by the slowest individual stage. To ensure that every part has enough time to complete its task, the system clock must match the duration of the most time-consuming step. This ensures synchronization across the entire line, even if other stages can finish their work more quickly, maintaining a steady flow of data.

167. (D) SSSD (Single Storage, Single Data)
Computers are categorized based on how many instructions and data streams they can process at the same time. Standard models include those that handle one of each, multiple of one, or multiple of both. These classifications provide a framework for understanding different architectural approaches to high-speed computing. Identifying valid categories is essential for analyzing the parallel processing capabilities of different hardware designs.

168. (B) P-3, Q-1, R-2
Pipelined processors face various challenges that can pause the execution flow. These include conflicts when two tasks need the same hardware, delays when data is not yet ready, and uncertainty caused by decision points in the program. Matching these specific hazards to their underlying causes is key for engineers who must design mechanisms like branch prediction and data forwarding to maintain efficiency.

NODIA

169. (B) It increases the instruction throughput of the processing unit.

Overlapping the various phases of command execution allows the processor to complete many more tasks in a given amount of time. While it does not necessarily make a single instruction finish faster, it significantly improves the overall volume of work performed. This increase in total output is the primary reason why this technique is used in almost all modern high-speed computer architectures.

170. (B) SIMD
Specialized processors that manage large collections of data through a single command fall into a specific architectural category. This design uses multiple processing elements that all perform the same task simultaneously but on different pieces of information. This approach is highly efficient for repetitive calculations across arrays, making it a cornerstone of high-performance scientific and graphics processing systems in contemporary computing.

171. (C) The time taken to execute a single instruction from start to finish (latency) is significantly reduced.
Dividing a task into stages improves the total number of jobs finished but does not inherently speed up the time it takes for any single job to move through the entire line. In fact, the added management and synchronization might slightly increase the total duration for one instruction. The primary benefit remains the volume of work completed rather than the speed of an individual task.

172. (B) An instruction needing the result of a preceding instruction that has not yet been computed.
Pipelined systems face delays when one step in a sequence requires information that is still being processed by an earlier step. This mismatch in timing creates a dependency that forces the system to pause until the necessary data is available. This common challenge requires sophisticated hardware solutions like data forwarding to minimize the impact on overall performance and maintain execution speed.

173. (A) Both A and R are true and R is the correct explanation of A.
Overlapping the different phases of multiple commands allows a processor to maintain a high rate of completion. By having several instructions in different stages of the process simultaneously, the hardware makes better use of all its components. Both statements correctly highlight this relationship, explaining why this technique is fundamental to achieving high throughput in modern, high-performance central processing unit designs.

174. (C) SIMD
Systems that apply a single command to an entire group of data elements follow a specific architectural model. This design is highly effective for tasks involving large arrays, where the same operation is repeated many times. By processing multiple pieces of information at once, the system achieves massive increases in speed for scientific and technical computations, distinguishing it from traditional sequential models.

175. (A) Both Statement I and Statement II are true.
Different types of commands are designed to handle varying amounts of information in a single step. One type operates on individual pieces of data, while another can process an entire collection of values simultaneously. Both statements accurately define these two categories of instructions, which represent different levels of architectural capability and are used in different types of computational tasks across various systems.

176. (B) (a), (b), and (c) only
High-performance computing for technical tasks often utilizes specialized hardware that processes large data collections through deeply structured stages. This approach is extremely efficient for complex mathematical models but does not rely on multiple independent processors or apply to simple text tasks. Understanding these core characteristics is essential for identifying the specific strengths and applications of advanced architectural designs used in modern scientific computing.

177. (C) The utilization of each pipeline stage over a sequence of clock cycles.
Visualizing the flow of tasks through a multi-stage processor involves mapping the activity of each part over time. This diagram shows exactly when each instruction enters and leaves different

Downloaded from www.pdf.tube

segments of the hardware. This is a vital tool for analyzing performance, identifying bottlenecks, and understanding how the system maintains high throughput by overlapping the execution of many different commands simultaneously.

178. (C) Array Processor - A synchronized set of processors executing the same instruction on different data.
Specialized parallel systems achieve high performance by using multiple identical units that work together under a single control. This organization allows the computer to process many pieces of data at once by applying the same logical task across the entire set. This synchronized approach is a key characteristic of certain high-speed architectures, particularly those used for large-scale mathematical and scientific simulations.

179. (B) Speed up the execution of complex arithmetic operations like floating-point multiplication.
Accelerating mathematical calculations often involves breaking down complex operations into smaller, consecutive steps. By using multiple hardware segments to handle different parts of the calculation at the same time, the processor can finish a long series of numeric tasks much more quickly. This specialized design is a vital component of high-performance systems that regularly perform heavy scientific or engineering arithmetic.

180. (A) Reduce the penalty incurred from control hazards.
Multi-stage processors face challenges when they encounter decision points in a program before the final outcome is known. To maintain speed, the hardware makes an educated guess about which path will be taken and continues processing. If the guess is correct, the system avoids a costly pause. This technique is essential for minimizing the delays caused by branching in complex software.

181. (B) External devices such as keyboards, monitors, and printers that communicate with the computer.
Computer systems rely on a variety of external hardware to interact with the outside world and users. These units provide the means to enter data, display results, and store information permanently. While they are not part of the core processing unit, they are essential for making the computer functional for practical tasks and allowing for communication between humans and the machine.

182. (A) To resolve differences between the CPU and peripherals, and supervise data transfer.
Interfacing hardware manages the complex communication between a central processor and various peripheral units. Key responsibilities include temporary data buffering, signal conversion between serial and parallel formats, and decoding unique device addresses. These circuits also generate necessary control signals to synchronize data transfers. By addressing these differences, the interface ensures smooth and reliable information flow throughout the entire computer system.

NODIA

183. (A) Both Statement I and Statement II are true.
Managing communication between components that do not share a common clock requires specific methods for synchronization. Some techniques use a single signal to trigger a transfer, while others use multiple signals to confirm that data has been successfully received. Both statements accurately describe these two fundamental approaches, which are used to ensure that information is moved reliably between different parts of the system.

184. (A) b, a, c, d
Direct data movement between storage and external hardware follows a specific sequence. First, the main processor sets the parameters, then it yields control of the communication paths. The specialized controller then performs the transfer and finally signals the processor that the task is finished. This organized flow allows for efficient bulk data movement with minimal direct intervention from the central unit.

185. (B) Interrupt-initiated I/O allows the CPU to execute other programs while waiting for I/O.
Efficient system design allows the central processor to avoid waiting for slower external hardware. By using signals that only alert the CPU when action is truly required, the system can focus on other computational tasks in the meantime. This multi-tasking capability is a cornerstone of modern operating systems,

ensuring that resources are used effectively while managing communication with peripheral devices.

186. (B) To allow a higher-priority device to interrupt the service of a lower-priority device.
Managing multiple requests for attention requires a system that can rank the importance of various tasks. This hierarchy ensures that critical events are handled immediately, even if the processor is already busy with less urgent activities. This ability to preempt tasks based on their importance is essential for maintaining the responsiveness and stability of complex, multi-device computer systems.

187. (D) Pipelined I/O
Different methods for moving information between the core processor and external hardware offer various levels of efficiency and speed. Common techniques include direct control, signal-based alerts, and specialized bulk transfers that bypass the main unit. While many architectural features use overlapping stages, this specific term is not a standard category for the fundamental modes of data transfer with peripherals.

188. (C) (a), (b) and (d) only
High-speed data movement is achieved by using a specialized controller that manages transfers without requiring the main processor to handle every individual word. This unit uses its own registers to track locations and amounts, significantly improving efficiency for large volumes of information. While the main unit can do other work, it is not necessarily idle, illustrating the benefit of delegating these tasks.

189. (C) To resolve the differences in speed, data format, and electrical characteristics.
Connecting a fast central processor to a variety of diverse external devices requires a bridge to manage their many differences. This bridge handles the timing, translates the physical signals, and adjusts the data formats so that the different parts can communicate effectively. This essential role ensures that information moves accurately and reliably between components that were not designed to work together directly.

190. (B) Manages data transfers between memory and I/O devices.
Specialized hardware can take over the task of moving large blocks of information, freeing the main processor for other work. This unit focuses exclusively on coordinating the flow of data between the computer's primary storage and various external units. By handling the details of these transfers independently, it significantly improves the overall efficiency and throughput of the entire computer system.

191. (D) Polling is the fastest method for identifying the highest-priority interrupt.
Methods for identifying which device needs attention vary in their efficiency. One technique requires the processor to manually check each potential source, which can be slow and wasteful. Other techniques use direct signals or physical connections that automatically identify the most important request. These alternative methods are generally much faster than manual checking, especially in systems with many different peripheral devices.

192. (A) P-3, Q-1, R-4, S-2
Managing the interaction between a processor and its environment involves various techniques for addressing, synchronization, and data movement. These include shared address spaces, signal-based confirmations, independent transfers, and manual status checking. Matching these terms to their definitions is essential for understanding how the computer coordinates the complex and varied tasks required for communicating with all its different external hardware components.

193. (B) Cycle stealing
High-speed data movement can occur by having a specialized controller use the system's communication lines during brief pauses in the main processor's activity. By taking control for only one or two cycles at a time, the controller can move information without significantly slowing down the overall execution of programs. This efficient method allows for parallel activity between data transfer and central processing.

194. (B) Serial communication
Transmitting information over long distances or with minimal wiring is often done by sending data

Click Here to Buy book on Amazon

as a sequence of individual bits. This approach uses a single path to move all parts of a character or word one after another. While it may be slower than using many parallel paths, its simplicity and lower hardware requirements make it a standard choice for many types of external connections.

195. (A) Both Statement I and Statement II are true.
Computers can be designed to interact with external hardware using two primary organizational styles. One treats peripheral units as distinct entities with their own instructions, while the other integrates them directly into the system's memory map. Both statements accurately describe these two methodologies, which represent different ways of organizing the communication between the central processor and the rest of the machine.

196. (A) Both A and R are true and R is the correct explanation of A.
Moving large amounts of information efficiently requires bypassing the central processor for routine data tasks. A specialized controller can handle the entire transfer directly between storage and external units, allowing the main processor to continue with more complex calculations. Both statements correctly highlight this relationship, explaining why this technique is essential for maintaining high system performance during heavy data movement tasks.

197. (C) Direct Memory Access (DMA)
Systems that require rapid response to high volumes of information need a transfer method that minimizes delay and processor overhead. By allowing data to move directly into storage without waiting for individual instructions from the central unit, the system can keep up with the fast pace of real-time signals. This method is the most effective choice for high-speed sensor data and other demanding hardware tasks.

198. (C) DMA - Reduces CPU overhead for bulk data transfer.
Different ways of moving information have unique characteristics that make them suitable for specific tasks. One technique is particularly valuable because it delegates the work of managing large transfers to specialized hardware, freeing the central processor to focus on more complex logical and mathematical jobs. Identifying this key benefit is essential for understanding the architecture of high-performance computer systems and their communication strategies.

199. (B) An exchange of control signals between the source and destination units to ensure synchronized data transfer.
Coordinating communication between different parts of a system requires a process of checking and confirming that both sides are ready. This involves sending signals back and forth to indicate when data is available and when it has been successfully received. This interaction ensures that information moves correctly and without loss, even when the two components operate at very different speeds or under different timing rules.

200. (C) (a), (b), (c), and (d) only
Interfacing hardware manages the complex communication between a central processor and various peripheral units. Key responsibilities include temporary data buffering, signal conversion between serial and parallel formats, and decoding unique device addresses. These circuits also generate necessary control signals to synchronize data transfers. However, they do not participate in the high-level application logic, which remains the processor's primary function in any standard computer architecture.

201. (B) Implementing a priority interrupt system.
Ranking the importance of different requests can be done by physically connecting devices in a sequential line. In this arrangement, the signal passes through each unit in order, allowing the one closest to the processor to take precedence. This simple and effective method ensures that multiple hardware requests are handled in a consistent hierarchy without requiring complex central arbitration logic or numerous dedicated wires.

202. (D) It is always faster than synchronous data transfer.
Moving data without a shared timing signal provides flexibility for connecting devices with very different speeds and characteristics. While this approach uses signals to coordinate each transfer, it does not inherently guarantee higher speed than methods that use a perfectly synchronized common clock. The primary

Downloaded from www.pdf.tube

advantage is the ability to communicate reliably across diverse hardware, rather than just increasing the absolute rate of information transfer.

203. (A) 960
Calculating the speed of information movement requires accounting for both the data bits and the necessary organizational bits, such as those that mark the start and end of a character. In a standard setup, each character requires a total of ten bits to be transmitted. By dividing the total bit rate by this value, we find the number of full characters that move through the line each second.

204. (B) To establish a hierarchy for handling simultaneous interrupt requests from multiple devices.
Managing multiple requests for attention requires a system that can rank the importance of various tasks. This hierarchy ensures that critical events are handled immediately, even if the processor is already busy with less urgent activities. This ability to preempt tasks based on their importance is essential for maintaining the responsiveness and stability of complex, multi-device computer systems in modern architectural designs.

205. (B) To indicate the beginning and end of a character for synchronization.
Transmitting data bit-by-bit requires a way for the receiver to know when each new piece of information starts and stops. By adding extra bits to each data packet, the sender provides the necessary timing cues for the two sides to stay in sync. This simple mechanism is essential for the reliable operation of asynchronous communication, allowing data to be correctly reconstructed at the destination.

206. (B) Main Memory
The primary storage unit of a computer holds the programs and data currently in use, providing relatively fast access for the central processor. It is directly addressable and is essential for the execution of every instruction the machine performs. This component acts as the central hub for information movement, connecting the high-speed processor with the slower, larger storage systems used for long-term retention.

207. (C) It is non-volatile and has a large storage capacity.
Large-scale storage systems are designed to keep information safe even when the computer is turned off. These devices can hold massive amounts of data compared to the primary memory, making them ideal for storing operating systems, applications, and user files permanently. While they are slower to access, their ability to retain information without power is a critical characteristic of all modern computing environments.

208. (C) Statement I is true, but Statement II is false.
Different types of storage are categorized by how they are searched and their place in the overall system. One specialized type finds information based on the data itself rather than its location, which is extremely fast for certain tasks. However, this is distinct from the high-speed buffer memory used to accelerate standard main storage. Both statements correctly highlight these organizational and functional differences in memory technology.

209. (B) (a), (b) and (c) only
Connecting small, fast buffers to larger, slower storage requires specific methods for organizing where data is placed. These techniques include mapping information directly to a single spot, allowing it to go anywhere, or using a hybrid approach that provides some choice. These common methods are essential for optimizing the speed of data retrieval and ensuring the processor can access information as quickly as possible.

210. (B) It allows the execution of programs that may be larger than the physical main memory.
Modern systems can run very large software by using part of the permanent storage to act as an extension of the primary memory. This approach gives the appearance of more space than physically exists, allowing the computer to manage multiple tasks and large files more flexibly. By moving pieces of data back and forth as needed, the system can handle complex jobs efficiently.

211. (A) 19 ns
Determining the average time to retrieve information involves combining the speeds of

the different storage layers and how often the needed data is found in the faster layer. In a typical setup, most requests are handled by the high-speed buffer, with only a few requiring the slower main storage. Calculating the weighted average of these two paths gives the true overall performance of the memory system.

212. (C) Physical location on the motherboard
Organizing storage systems into levels is based on the trade-offs between cost, capacity, and speed. Layers that are faster are also more expensive and smaller, while layers that hold more data are cheaper but slower to access. This hierarchy is a logical structure designed to optimize performance, and it does not depend on the physical placement of the components on the circuit board.

213. (C) a, b, c, d
Storage components in a computer follow a clear progression of speed, from the extremely fast units inside the processor to the much slower permanent storage devices. Understanding this order is essential for analyzing system performance, as each level serves as a buffer for the one below it. This structured organization ensures that the most frequently used information is always accessible at the highest possible speed.

214. (B) To act as a buffer between the CPU and main memory, holding frequently accessed data.
High-speed buffer memory is used to bridge the significant speed gap between the central processor and the much slower primary storage. By keeping copies of recently used information in a smaller, faster location, the system can provide the data the processor needs with much less delay. This significantly improves the overall speed of program execution and the efficiency of the entire computer system.

215. (B) To translate logical addresses generated by the CPU into physical addresses in main memory.
Managing how software interacts with physical storage requires a specialized unit that maps the abstract locations used by programs to real locations in the memory hardware. This translation allows the system to protect different programs from each other and manage memory more flexibly. This role is fundamental to the operation of modern operating systems that handle multiple tasks and use sophisticated storage techniques.

216. (D) The entire program must be loaded from disk into memory before it can begin execution.
Advanced memory management allows a computer to start running a program even if only part of it is present in the primary storage. This on-demand approach is far more efficient than requiring everything to be loaded first, as it saves space and allows the machine to respond more quickly. Additional pieces of the program are brought in from permanent storage only as they are needed.

217. (C) The CPU tends to access memory locations that are near other recently accessed locations or are accessed repeatedly.
Modern computing relies on the fact that programs do not access information randomly. Instead, they often work on small chunks of data or repeat certain sections of code many times. This predictable behavior allows high-speed buffers to be extremely effective, as they can accurately guess which information will be needed next. This principle is why relatively small amounts of fast memory can significantly accelerate overall performance.

218. (C) (a), (b), (c) and (d) only
Storage levels are structured such that moving down the hierarchy increases capacity and access time while decreasing the cost per bit. Moving up the hierarchy puts the data closer to the processor, where it is accessed much more frequently. While most levels have these characteristics, they are not all volatile, as permanent storage is designed to keep data even when power is lost.

219. (C) Cache Memory
Managing data consistency between high-speed buffers and the primary storage requires specific rules for when and how changes are recorded. Some methods update both locations at the same time, while others wait until a piece of data is being removed from the buffer before saving the changes to the main memory. These policies are critical for ensuring that the most recent information is preserved.

220. (A) Both Statement I and Statement II are true.

Click Here to Buy book on Amazon

Semiconductor storage comes in two main varieties with different performance characteristics. One type is extremely fast but takes up more space and is used for high-speed buffers. The other is denser and cheaper but requires regular power pulses to keep its data from fading. Both statements accurately describe these technologies, which are the fundamental building blocks of all modern computer memory systems.

221. (C) Set 5
Mapping data between different storage levels involves a mathematical calculation to determine the destination. In a hybrid arrangement that provides some choice of placement, the specific location is found by taking the block number and applying a remainder operation based on the total number of sets. This ensures that information is distributed across the faster storage in a predictable and organized way for retrieval.

222. (D) Cache (SRAM) - Nanoseconds
Different types of storage are characterized by the time it takes to retrieve a piece of information. The fastest layers, located closest to the processor, operate in mere billionths of a second. As you move away from the core, these times increase significantly for main storage and even more for mechanical storage devices. Correctly matching these speed levels is essential for understanding the overall memory hierarchy.

223. (B) A data structure in memory that maps virtual page numbers to physical frame numbers.
Managing the transition between abstract program locations and physical hardware requires a detailed reference guide. This map allows the system to keep track of exactly where each chunk of a program's data is stored in the real memory chips. By looking up these entries, the processor can correctly find and use the information it needs, even if it is scattered across different physical areas.

224. (B) Parallel search by content.
Some specialized storage units are designed to find information by looking for a specific value throughout the entire memory at the same time. This is much faster than checking each location one by one. This unique capability is essential for certain high-speed tasks, such as managing translations and lookups in network hardware or advanced memory management units within a central processor.

225. (D) Virtual memory is always faster to access than physical main memory.
Using sophisticated storage techniques provides many benefits, such as the ability to run larger programs and protect data. However, these techniques often introduce a small amount of extra work for the hardware and software, which can result in slightly longer access times than direct physical memory. The trade-off is necessary for the advanced features and flexibility that modern computing systems require to operate effectively.

226. (A) Index into the page table to find the corresponding frame number.
Determining the final physical location of data in a complex system requires a structured lookup process. The hardware uses one part of the address as a key to search a comprehensive reference guide. This guide reveals exactly which section of the physical memory chips holds the required information. This multi-step translation is a core part of how modern processors manage large-scale data and software tasks.

227. (D) The system is spending an excessive amount of time swapping pages between main memory and auxiliary memory.
Efficient operation can break down if a computer tries to run too many programs at once or if the programs need more fast memory than is available. When this happens, the system spends all its resources moving pieces of data back and forth between primary and permanent storage, leaving no time for actual work. This condition significantly degrades performance and can make the computer unresponsive.

228. (A) Least Recently Used (LRU)
To make room for new data in a limited high-speed buffer, the system must decide which existing piece to remove. A common and effective strategy is to discard the information that has not been accessed for the longest period of time. This approach assumes that data that hasn't been used recently is the least likely to be needed again soon, maximizing the hit rate.

NODIA

229. (B) An event where the CPU finds the required data in the cache.
The ultimate goal of a multi-level storage system is to have the processor find the information it needs in the fastest possible location. When this happens, the system avoids a costly wait for the slower primary memory. This successful retrieval is a key metric for measuring the efficiency of the memory hierarchy and the overall performance of the computer while running various software.

230. (B) Magnetic tapes and optical disks are examples of auxiliary memory devices.
Storage systems that provide massive capacity and keep data safe without power are essential for modern computers. These include both traditional mechanical units and newer electronic designs that hold everything from system files to personal photos. These devices are not used for immediate instruction execution but serve as the vast, permanent repository for all the information and software used by the computer system.

231. (C) It contains two or more CPUs that share access to a common memory and peripherals.
Increasing computational power is often achieved by building systems with multiple independent processing units that work together. These units are connected so they can all use the same primary storage and external hardware. This organization allows the computer to perform many complex tasks simultaneously, significantly improving throughput and reliability for demanding server and high-performance workstation applications in modern environments.

232. (A) A crossbar switch
Connecting multiple processors to multiple memory modules requires a sophisticated internal communication structure. One highly effective design provides a dedicated path between every possible pair of components, allowing many simultaneous transfers without any conflict. While this is more expensive to build than a shared bus, it offers superior performance for high-end systems that require the fastest possible movement of information between units.

233. (A) Both Statement I and Statement II are true.
Systems with multiple processors can be organized in two primary ways based on how they share resources and communicate. One style uses a single shared memory and operating system for close coordination, while the other connects independent machines through a network. Both statements accurately describe these fundamental architectural choices, which dictate how the overall system scales and how it handles complex distributed tasks.

234. (B) (a), (b) and (c) only
Managing shared resources in a multi-processor system requires rules for deciding which unit gets access when multiple requests occur simultaneously. These can include physical sequential connections, parallel logic for faster decisions, or dynamic algorithms that rotate access or prioritize certain units. While semaphores are used for software-level coordination, these hardware mechanisms are essential for the physical management of the system's communication buses.

235. (B) It is necessary to coordinate activities and ensure correct data sharing among processors in a multiprocessor system.
When multiple units work together on a single problem, they must have ways to share results and ensure they do not interfere with each other. This requires constant communication and the use of special signals to synchronize their actions. Without these mechanisms, the different processors could overwrite each other's work or become stuck waiting for data, leading to incorrect results or system failures.

236. (A) The cache coherence problem.
Systems with multiple independent buffers face a challenge when one processor changes a value that others are also using. If the other processors continue to use their own local copies, they will have outdated information, leading to errors. Maintaining consistency across all these separate storage areas is a critical requirement for any multi-processor architecture to ensure the accuracy and reliability of shared data.

237. (D) Single-port register file
Connecting processors and memory requires specialized structures that allow for efficient and often simultaneous data movement. These include shared buses, complex switch networks, and memory units with multiple access points. While internal registers are

Downloaded from www.pdf.tube

essential for individual units, they do not serve as the primary infrastructure for connecting the major components of a multi-processor system together for large-scale information sharing and communication.

238. (B) A chip containing two or more independent processing units known as cores.
Technological advancements have allowed engineers to place multiple complete processing engines onto a single piece of silicon. Each of these units can execute its own tasks and share internal resources like high-speed buffers. This design provides a massive increase in performance while keeping the physical size and power requirements low, making it the standard architecture for modern computers from laptops to servers.

239. (A) Ensuring that all caches have the most recent and correct value of a shared data item.
Maintaining consistency across multiple high-speed buffers is vital when several processors are working on the same set of information. If one unit modifies a value, the system must ensure that all other units are aware of the change and do not use old, incorrect data. This coordination is essential for the logical integrity and mathematical accuracy of all multi-processor computing tasks and systems.

240. (B) The daisy-chaining method is a form of parallel arbitration.
Deciding which processor can use a shared resource involves various hardware techniques with different performance characteristics. One common method passes a signal along a sequence of units, which is a serial rather than parallel process. Other methods use fixed or rotating rules to give all units a fair or prioritized chance at access. Understanding these distinctions is key for analyzing how multi-processor systems manage internal communication.

241. (A) The need to prevent multiple processors from simultaneously modifying shared data, which could lead to inconsistent results.
Coordinating access to shared information is critical when multiple tasks are running at the same time. Special software tools act as digital locks, ensuring that only one process can change a specific piece of data at any given moment. This prevents the race conditions that occur when simultaneous updates cause data corruption, ensuring that the final results of calculations are always predictable and correct.

242. (B) (a), (b), and (c) only
Advanced computing systems aim to improve total output and reliability by using many units that can share the workload. These systems require complex mechanisms for communication and synchronization but are not limited to simple bus designs or basic sequential classifications. Identifying these core properties is essential for understanding the true capabilities and challenges of building and using large-scale, high-performance multi-processor machines in modern computing environments.

243. (B) Maintain cache coherence in bus-based multiprocessor systems.
Consistency across multiple storage buffers can be maintained by having each unit constantly listen to the activity on the shared communication path. If a unit sees that a value it holds is being updated by another processor, it can immediately mark its own copy as outdated. This automatic monitoring is a highly effective way to ensure that all parts of the system are always working with the most current data.

244. (A) Both Statement I and Statement II are true.
Multi-processor designs vary in how they organize the connection to primary storage. One approach provides identical access for all units, while another creates a more complex structure where some data is closer and faster to reach than other data. Both statements accurately describe these two common architectural styles, which have a major impact on how software must be designed to achieve the best performance.

245. (C) Hypercube - A multidimensional binary cube structure.
Connecting a vast number of processors can be done using various logical shapes that provide efficient paths between any two units. One advanced design uses a structure where each node is connected to many others in a pattern that scales well as more units are added. Correctly matching these complex topologies to their names is essential for understanding the

NODIA

infrastructure of modern supercomputers and large-scale parallel systems.

246. (A) The mechanism for deciding which processor gets access to a shared resource when there are multiple simultaneous requests.
Managing a shared communication path in a multi-unit system requires a way to resolve conflicts when more than one unit wants to use it at once. This hardware-based decision process ensures that access is granted in an organized way, preventing data collisions and ensuring that the system continues to function smoothly. It is a fundamental component of the infrastructure that allows multiple processors to work together.

247. (C) Managing communication, synchronization, and data consistency among the many processors.
The primary difficulty in building extremely large systems is not the physical hardware itself, but the immense logical complexity of coordinating all the different units. Ensuring that they all have the correct data and work together in perfect harmony requires sophisticated software and hardware design. Overcoming these challenges is the key to creating systems that can solve the world's most complex scientific and engineering problems.

248. (B) Crossbar switch
An interconnection structure that provides a separate path from each memory module to each processor, eliminating bus contention but with high hardware cost, is the crossbar switch. This architecture allows multiple simultaneous data transfers without any interference, offering the highest possible performance. However, its complexity and cost increase significantly as more processors and memory modules are added, making it a specialized choice for high-end systems.

249. (A) They provide increased performance for parallel tasks within a smaller power and thermal envelope.
Integrating multiple processing units onto a single piece of silicon offers massive benefits for modern computing. By sharing internal resources and reducing the physical distance between components, these designs can perform more work while using less energy and generating less heat than older, multi-chip systems. This efficiency is why this architecture is used in everything from mobile phones to the most powerful data center servers.

250. (B) Controlling access to a shared resource by multiple processes to achieve synchronization.
Multi-processing environments require robust mechanisms to manage shared hardware or data access. These variables act as signaling tools to prevent concurrent modification of critical sections, which could cause system instability or data corruption. By implementing atomic operations for incrementing or decrementing values, they ensure that only authorized processes utilize resources at any given time, maintaining overall system architectural and data integrity.

251. (B) To reduce the complexity and resource requirements of the circuit.
Boolean Algebra and Map Simplification techniques provide a mathematical and graphical framework for minimizing logical expressions. Reducing the number of required logic gates directly decreases production costs, power consumption, and signal propagation delays. Efficient circuit design relies on these methods to ensure that digital systems use the minimum amount of hardware resources.

252. (A) Combinational circuits lack memory, whereas sequential circuits incorporate it.
Logic circuits are distinguished by their functional dependency on time and memory components. Combinational circuits generate outputs solely from their current inputs without any internal storage. In contrast, sequential circuits utilize memory elements like flip-flops to retain past information. This allows their outputs to depend on both current inputs and the historical sequence of previous inputs.

253. (C) systematically minimizing a logical expression.
The Karnaugh Map offers a visual approach to simplifying Boolean functions without performing extensive algebraic manipulations. By grouping adjacent ones in a grid representation, designers can identify redundant terms in a logical expression. This systematic reduction process is essential for creating streamlined digital circuits that require fewer gates and perform logic operations with greater efficiency.

Click Here to Buy book on Amazon

254. (D) A combinational circuit, as its output depends only on the present input.
Decoders and multiplexers are standard examples of logic structures where the output state is determined entirely by the binary values present at the inputs at that specific moment. These devices do not possess internal memory components to store previous states. Therefore, they fall under the classification of combinational circuits rather than sequential circuits or memory units.

255. (C) The rules and theorems of Boolean logic.
The entire architecture of digital computing is built upon the foundational layer of logic gates. The behavior, interconnection, and optimization of these gates are governed by the rules and theorems of Boolean Algebra. This logical framework provides the necessary mathematical tools for the analysis and synthesis of everything from basic circuits to highly complex integrated systems.

256. (B) It simplifies the hardware for subtraction by converting it to addition.
Using 2's complement representation allows digital hardware to perform subtraction using the same adder circuits used for addition. This is achieved by adding the 2's complement of the subtrahend to the minuend. Additionally, this system provides a unique representation for zero, eliminating the ambiguity found in the 1's complement system, which simplifies arithmetic logic.

257. (D) Floating-point representation, due to its wide dynamic range.
Scientific calculations often involve values that span many orders of magnitude. Floating-point representation uses a mantissa and an exponent to accommodate this extremely wide dynamic range, allowing for the representation of both very large and very small numbers. While fixed-point is faster for specific ranges, it lacks the flexibility required for complex engineering and scientific computational tasks.

258. (A) They ensure the integrity of data by flagging potential corruption.
Error detection codes are implemented to maintain data reliability during storage and transmission. For instance, a parity bit added to a binary word can identify if a single-bit error has occurred. By detecting these discrepancies, the system can flag corrupted information. These codes serve as a critical safeguard for data integrity across various digital communication and storage interfaces.

259. (C) Fixed-point has an implied radix point, while floating-point uses an explicit exponent.
Fixed-point representation assumes the radix point is at a consistent, predefined position within the numeric string, making it suitable for hardware with limited range requirements. Floating-point representation, however, uses an explicit exponent field to dynamically shift the radix point. This distinction allows floating-point systems to handle a significantly larger variety of real numbers compared to fixed-point formats.

260. (A) they offer different ways to represent the same underlying binary data.
While digital systems natively use binary for internal processing, octal and hexadecimal systems provide more compact and human-readable representations of the same binary values. Converting between these bases is essential for tasks like debugging or programming, as it allows for easier visualization of large bit strings. All three systems describe the same underlying data in different numerical formats.

261. (D) In the same memory as the data they operate on.
Stored program organization is the central principle of modern computer architecture. This concept dictates that both the executable instructions of a program and the data required for those instructions are housed together in the same main memory. This unified storage approach allows the CPU to fetch both instructions and operands using a standardized memory access mechanism during execution.

262. (B) Fetch instruction using address in PC, place instruction in IR, decode instruction.
The instruction cycle starts with the fetch phase, where the address stored in the Program Counter is used to retrieve an instruction from memory. This instruction is then loaded into the Instruction Register. Subsequently, the Timing and Control unit decodes the instruction's opcode. This sequential process ensures that the

NODIA

processor correctly identifies and prepares for the required operation.

263. (C) generate the signals required to direct the steps of the instruction cycle.
The Timing and Control unit serves as the coordinator of the processor's internal activities. After decoding the opcode of an instruction, it generates the specific control signals needed to manage data flow and execution. These signals synchronize the actions of the memory, registers, and arithmetic units, ensuring that each step of the instruction cycle is performed in order.

264. (D) pause the current sequence of operations to handle a specific request.
An interrupt serves as a signaling mechanism that allows external devices or internal conditions to gain the immediate attention of the CPU. When an interrupt occurs, the processor suspends its current task, saves its state, and executes a specific routine to handle the request. Once the event is processed, the computer resumes the previous program from the point of interruption.

265. (A) To hold the instruction that is currently being decoded and executed.
During the fetch phase of the instruction cycle, the retrieved binary instruction is placed into the Instruction Register. This specific register maintains the instruction while the control unit decodes the operation code and determines the necessary execution steps. The Instruction Register acts as a temporary buffer that holds the current instruction until its processing cycle is fully completed.

266. (B) The presence of a Control Memory.
A microprogrammed control unit is distinguished by its use of Control Memory, which stores microprograms. Unlike hardwired units that use a fixed arrangement of logic gates to generate signals, a microprogrammed unit retrieves microinstructions from this internal memory. This architectural feature allows the control logic to be defined by software-like sequences rather than permanent hardware configurations of logic circuits.

267. (B) selecting the address of the next microinstruction to be executed.
Address sequencing is the process used within a microprogrammed control unit to determine which microinstruction should be executed next. The address can be derived from the machine instruction's opcode, status flags, or branch information within the current microinstruction. This mechanism provides the flexibility needed to implement complex logical branching and subroutines directly within the control logic of the processor.

268. (A) the relative ease of implementing complex instructions and modifying control logic.
Microprogrammed control units offer significant flexibility compared to hardwired designs. Because the control signals are generated by microprograms stored in memory, designers can implement complex instruction sets more systematically. Modifying or correcting the control logic simply requires updating the microinstructions in the Control Memory. This flexibility makes it easier to support sophisticated architectures without redesigning the physical hardware.

269. (C) the execution of a sequence of microinstructions.
In a microprogrammed system, the execution of a machine-level instruction triggers a specific microprogram. This microprogram consists of a sequence of microinstructions, each of which produces elementary control signals. These signals drive the internal operations of the CPU. Consequently, the control unit's behavior is the direct result of executing these stored microinstruction sequences rather than hardwired logical paths.

270. (D) A set of microinstructions stored in Control Memory to implement one machine instruction.
A microprogram is a series of low-level microinstructions designed to perform the individual steps required to execute a single machine language instruction. These microprograms are housed within the Control Memory. When a machine instruction is decoded, the control unit executes the corresponding microprogram. This hierarchical approach allows for complex operations to be broken down into simpler, manageable hardware control steps.

NODIA

271. (C) It offers greater flexibility by allowing multiple registers for operands and results.
General Register Organization provides the CPU with a set of versatile registers that can store various operands and intermediate results. This design is more flexible than an accumulator-based architecture, which is restricted to using a single dedicated register for arithmetic operations. By using multiple registers, the processor can handle complex expressions and data manipulations more efficiently during program execution.

272. (B) Implicitly at the top of the stack structure.
Stack Organization manages data using a last-in, first-out sequence. In this architecture, arithmetic instructions do not require explicit operand addresses because the values are implicitly located at the top of the stack. The processor automatically pops the required operands to perform the operation and then pushes the result back onto the stack, resulting in shorter and simpler instruction formats.

273. (D) specify how the location of an operand is determined.
Addressing Modes define the specific rules used by the CPU to locate operands required for instruction execution. These modes allow for various ways of accessing data, such as finding it directly within the instruction, at a specific memory address, or through a pointer. This variety provides programmers with the flexibility to access memory efficiently using different data structures and offsets.

274. (A) Immediate addressing.
Immediate addressing is a mode where the operand itself is included as a part of the instruction bit string. Unlike other modes that require fetching data from a register or a memory location, this method allows the CPU to access the value immediately upon decoding the instruction. This is highly efficient for using constants or initial values in a program.

275. (C) To structure the bits of an instruction to represent opcode, operands, and mode.
Instruction Formats define the organizational layout of bits within a machine instruction. They specify which bits represent the operation to be performed, which bits identify the operands, and which bits indicate the addressing mode. By establishing these structures, the format ensures that the control unit can correctly decode and execute each instruction according to the processor's programming model.

276. (A) To increase the overall instruction throughput of the processor.
An instruction pipeline enhances performance by overlapping the execution of multiple instructions. While a single instruction's total processing time remains the same, the pipeline allows a new instruction to finish in nearly every clock cycle. This concurrent processing significantly increases the total number of instructions the CPU can complete over time, which is known as instruction throughput.

277. (D) Parallel processing.
Pipelining is a fundamental implementation of parallel processing, where several operations are executed simultaneously. By dividing the instruction cycle into distinct stages that operate concurrently on different instructions, the system achieves temporal parallelism. This strategy is essential in modern computer architecture to maximize resource utilization and improve the overall computational efficiency of the central processing unit and arithmetic units.

278. (B) accelerate complex numerical calculations by breaking them into stages.
Arithmetic pipelines are designed to speed up complex mathematical operations, such as floating-point addition or multiplication. These operations are decomposed into smaller, sequential sub-operations that can be executed in parallel across different pipeline stages. This allows multiple calculations to be in progress simultaneously, resulting in a much faster output rate for intensive numerical tasks in scientific and engineering applications.

279. (C) operations on large, one-dimensional arrays of data.
Vector processing is highly efficient for applications that perform identical operations on large, one-dimensional arrays, or vectors. By using specialized hardware to process multiple data elements simultaneously, it exploits data parallelism. This approach is particularly beneficial for scientific simulations and engineering tasks where repetitive calculations

Downloaded from www.pdf.tube

on datasets are common, providing a significant performance boost over traditional scalar processing methods.

280. (D) Multiple instructions are in different stages of execution at the same time.
In an instruction pipeline, the total execution cycle is split into several independent stages, such as fetch, decode, and execute. As one instruction moves to the next stage, a new instruction enters the pipeline. This ensures that different parts of the CPU hardware are working on multiple instructions concurrently, thereby optimizing efficiency and increasing the rate of instruction completion.

281. (A) To resolve the differences between the CPU and peripheral devices.
An Input-Output Interface acts as a mediator between the central processor and external peripherals. It is responsible for bridging differences in operational speeds, data formats, and control signal requirements. By normalizing these variations, the interface allows the CPU to communicate effectively with diverse devices such as keyboards, printers, and disks, ensuring seamless data exchange within the system.

282. (B) Programmed I/O.
In programmed I/O mode, the CPU is responsible for managing the entire data transfer process. It must continuously poll the peripheral device to check its status and then execute the transfer instructions itself. This method keeps the processor fully occupied during the duration of the I/O operation, preventing it from performing other useful tasks until the transfer is complete.

283. (B) It provides a highly efficient method for transferring large data blocks without CPU intervention.
Direct Memory Access allows a dedicated controller to manage the transfer of large blocks of data between a peripheral and main memory. Unlike interrupt-driven I/O, which requires the CPU to handle each data word, DMA takes control of the system bus and performs the transfer autonomously. This significantly reduces the overhead on the CPU, allowing it to execute other programs.

284. (A) coordinate the exchange of data when a shared clock is not used.
Asynchronous data transfer is used when the sender and receiver do not share a common clock signal. To ensure data is transferred correctly, the devices use handshaking signals like strobe and acknowledge. These signals coordinate the timing of the exchange, indicating when data is valid and when it has been successfully received, thus preventing data loss during the communication process.

285. (C) multiple devices may request the CPU's attention simultaneously.
A priority interrupt system is essential for managing situations where several peripheral devices generate interrupt requests at the same time. The system evaluates the importance of each request based on predefined priority levels. It then ensures that the CPU services the most critical or highest-priority device first, maintaining an orderly and efficient response to multiple competing input-output events.

NODIA

286. (D) To bridge the speed mismatch between the CPU and Main Memory.
Cache memory is a small, high-speed storage unit located near the CPU. Its primary role is to store frequently used instructions and data, reducing the time the processor must wait for slower main memory accesses. By bridging this performance gap, cache memory significantly improves overall system speed, allowing the fast CPU to operate closer to its maximum potential frequency.

287. (C) allow programs to use a larger address space than the physical main memory.
Virtual memory is a management technique that creates the illusion of a very large main memory for applications. It allows programs to use an address space that exceeds the physical capacity of the DRAM by dynamically swapping data between main memory and auxiliary storage. This enables the execution of larger programs and improves the efficiency of multitasking environments.

288. (B) The principle of locality of reference.
The effectiveness of cache memory is based on the principle of locality of reference. This principle observes that programs tend to access data and instructions that are spatially or

temporally close to those recently accessed. By keeping these relevant items in the high-speed cache, the system increases the probability of fast data retrieval, thereby optimizing the processor's execution speed.

289. (D) It allows for data retrieval based on content rather than address.
Associative memory is a specialized type of hardware storage where data is accessed by matching its content rather than by using a numerical memory address. This capability allows for extremely rapid search operations, as all memory locations can be compared simultaneously. It is particularly useful in memory management units for quickly searching page tables and in cache systems for tag comparisons.

290. (A) Memory Management Hardware.
In a virtual memory system, the program uses logical addresses that do not correspond directly to physical locations in DRAM. The Memory Management Hardware, or MMU, is responsible for performing the real-time translation of these logical addresses into physical addresses. This hardware ensures that data is correctly retrieved from or stored in the actual memory locations assigned by the system.

291. (C) It contains multiple CPUs that share access to a common memory.
A multiprocessor system is defined by the presence of two or more central processing units that operate simultaneously. These processors are interconnected and share access to a common main memory and peripheral resources. This architecture enables the concurrent execution of multiple instruction streams, which significantly boosts the system's throughput and allows for more powerful parallel computing capabilities.

292. (A) maintaining data consistency across multiple private caches.
Cache coherence is a critical issue in shared-memory multiprocessors where each CPU has its own local cache. If multiple processors store a copy of the same memory location and one modifies it, the other copies become stale. Protocols are necessary to ensure that all processors see the most recent and consistent version of data, preventing errors during parallel execution.

293. (D) To facilitate communication paths between processors and memory.
Interconnection structures, such as crossbar switches or common buses, provide the physical and logical paths through which processors communicate with each other and access shared memory. A crossbar switch, in particular, allows multiple processors to connect to different memory modules simultaneously. These structures are essential for managing data flow and ensuring efficient resource utilization in a multiprocessor environment.

294. (B) resolve conflicts when multiple processors request a shared resource.
Interprocessor Arbitration is a vital mechanism in multiprocessor systems that manages access to shared resources like the system bus or main memory. When multiple CPUs attempt to use the same resource at the same time, the arbitration logic determines which processor should be granted access. This prevents data collisions and ensures that shared resources are allocated in a fair, orderly manner.

295. (C) an implementation of multiprocessor architecture on a single chip.
Multicore processors represent the integration of multiple independent processor cores onto a single integrated circuit. This design brings the concepts of multiprocessor architecture to a smaller scale, allowing for parallel execution within a single chip. Each core can handle its own instruction stream while sharing access to memory and caches, making advanced parallel computing accessible in modern consumer devices.

296. (D) Complex instructions that can perform multi-step operations.
Complex Instruction Set Computer architecture is defined by its extensive instruction set, where a single instruction can perform multiple operations. These can include complex tasks like simultaneously multiplying two numbers and storing the result in memory. By narrowing the semantic gap between high-level languages and machine code, CISC simplifies the compiler's task, though it often requires microprogrammed control.

Click Here to Buy book on Amazon

297. (A) having the compiler combine a sequence of simple instructions.
The Reduced Instruction Set Computer philosophy relies on a small set of simple, optimized instructions. Instead of having single instructions for complex tasks, the RISC architecture uses the compiler to synthesize these operations from a sequence of basic instructions. This approach allows for faster hardware execution and simplified processor design, as each individual instruction performs only a very specific task.

298. (B) Because the uniformity and single-cycle nature of their instructions simplify pipeline flow.
RISC processors are highly efficient for pipelining because their instructions have fixed lengths and generally execute in a single clock cycle. This uniformity reduces the complexity of the pipeline stages and minimizes the occurrence of pipeline stalls. By maintaining a steady flow of instructions, RISC architectures achieve higher throughput and better utilize the overlapping execution stages compared to complex architectures.

299. (B) all data must be loaded from memory into registers before being operated on.
A load/store architecture, common in RISC designs, separates memory access from data processing. Only specific load and store instructions are permitted to interact with main memory. All arithmetic and logical operations must be performed on data already present in the processor's general-purpose registers. This simplification streamlines the instruction set and makes the execution of instructions more predictable.

300. (A) A microprogrammed unit, because it provides the flexibility to manage a complex instruction set.
CISC processors typically utilize microprogrammed control units to handle their large and varied instruction sets. Because CISC instructions can be highly complex and variable in length, implementing them with hardwired logic would be extremely difficult. Storing the control logic as microprograms in memory provides the necessary flexibility to define and modify these intricate instructions systematically during the design phase.

www.nodia.in

NODIA

operators for user-defined types.
Statement II : All operators in C++ can be overloaded without any restrictions.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is true, but Statement II is false.

37. A homogeneous coordinate representation of a 2D point (x, y) is given by:
(A) (x, y, 0)
(B) (x/w, y/w)
(C) (x, y, 1)
(D) (x, y)

38. Which of the following is a key advantage of using composite transforms in computer graphics?
(A) They reduce the number of vertices in a polygon.
(B) They combine a sequence of transformations into a single matrix, improving efficiency.
(C) They allow individual transformations to be applied more slowly.
(D) They convert 2D transformations into 3D transformations.

39. Consider the following statements regarding properties of scalar data types:
(a) It is composed of other data types.
(b) It is atomic and cannot be further decomposed.
(c) Integers and floating-point numbers are examples.
(d) An array is a primary example of a scalar type.
Which of the above statements are true?
(A) (a) and (d) only
(B) (b) and (c) only
(C) (a), (c) and (d) only
(D) (b), (c) and (d) only

40. In C programming, file handling functions like `fopen()`, `fclose()`, and `fread()` are found in which header file?
(A) `stdio.h`
(B) `string.h`
(C) `stdlib.h`
(D) `math.h`

41. What is the primary purpose of the `argv` parameter in the `main` function in C (i.e., `int main(int argc, char *argv[])`)?
(A) To hold the number of arguments passed to the program.
(B) To specify the return value of the program.
(C) To store an array of strings representing the command line arguments.
(D) To count the number of variables used in the program.

42. The concept of Subprogram Control in a programming language like C primarily deals with:
(A) The mechanisms for passing parameters to and returning values from functions.
(B) The flow of execution within a single function.
(C) The management of loops and conditional statements.
(D) The control of low-level hardware interrupts.

43. Which OOP concept is best described as providing only essential information to the outside world and hiding the background details?
(A) Inheritance
(B) Polymorphism
(C) Instantiation
(D) Abstraction

44. Consider the following statements regarding features of Web Programming technologies:
(a) Java Applets are small Java programs that run within a web browser.
(b) XML's primary purpose is to define the presentation and styling of a document.
(c) Servlets are Java programs that run on the web server, extending its capabilities.
(d) DHTML is a specific version of HTML, standardized by W3C as HTML 5.
(e) Scripting languages like JavaScript are typically interpreted rather than compiled.
Which of the above combinations is correct?
(A) (a), (b), and (d) only
(B) (b), (d), and (e) only
(C) (a), (c), and (e) only
(D) (c), (d), and (e) only

Downloaded from www.pdf.tube

9. Match the following C language concepts in List I with their descriptions in List II.

List I: (Concept)

P. Pointer

Q. Structure

R. Union

S. Preprocessor Directive

List II: (Description)

1. A user-defined data type that can hold elements of different types in a shared memory location.
2. A variable that stores the memory address of another variable.
3. An instruction to the compiler that is processed before actual compilation.
4. A user-defined data type that groups together variables of different data types.

(A) P-2, Q-4, R-1, S-3
(B) P-1, Q-2, R-3, S-4
(C) P-2, Q-1, R-4, S-3
(D) P-4, Q-2, R-1, S-3

10. Which of the following is an example of a composite data type?
(A) array
(B) integer
(C) float
(D) boolean

11. What is the role of a virtual computer in the context of language design?
(A) A physical hardware machine for executing software.
(B) A set of hardware-independent development tools.
(C) A tool to translate code into machine instructions.
(D) A hardware abstraction providing a language interface.

12. Consider the following statements regarding C pointers:

Statement I : A pointer variable stores the value of the variable it points to.

Statement II : The '&' operator is used to get the address of a variable, and the '*' operator is used to dereference a pointer.

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

13. In 2-D transformations, what does a shearing transformation do?
(A) It moves an object along a straight-line path.
(B) It alters the size of an object.
(C) It distorts the shape of an object by sliding layers over each other.
(D) It rotates an object about the origin.

14. Consider the following statements about programming language paradigms:
(a) The procedural paradigm is based on the concept of procedure calls.
(b) The object-oriented paradigm is organized around objects rather than actions.
(c) The functional paradigm treats computation as the evaluation of mathematical functions.
(d) The logic paradigm is based on formal logic and rule-based queries.

Which of the above statements are correct?
(A) (a) and (b) only
(B) (a), (c) and (d) only
(C) (b) and (d) only
(D) (a), (b), (c) and (d)

NODIA

15. What does XML stand for?
(A) Extensible Markup Language
(B) Extended Markup Language
(C) Example Markup Language
(D) Executable Markup Language

16. Which of the following is NOT a fundamental property of a data type?
(A) The set of values the object can take.
(B) The memory address of the object.
(C) The set of operations applicable to the object.
(D) The way values are represented in memory.

17. In C programming, what is the primary purpose of the `union` keyword?
(A) To group multiple variables of different types under a single name.
(B) To allow different variables to share the same memory location.
(C) To define a new data type that is an alias for an existing type.
(D) To create a constant variable whose value cannot be changed.

18. In the context of C++, an abstract class is a class that:

Q. Rotation by angle θ
R. Reflection about y-axis
S. Shear in x-direction by shx
List II: (Matrix Representation)
1. `[[cosθ, -sinθ], [sinθ, cosθ]]`
2. `[[-1, 0], [0, 1]]`
3. `[[Sx, 0], [0, Sy]]`
4. `[[1, shx], [0, 1]]`
(A) P-3, Q-1, R-2, S-4
(B) P-1, Q-3, R-4, S-2
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-2, S-3

55. In C++, what is the difference between a class and an object?
(A) A class is an instance of an object.
(B) An object is a blueprint, while a class is an instance.
(C) There is no difference; the terms are interchangeable.
(D) A class is a blueprint, while an object is an instance created from that blueprint.

56. Which of the following is NOT a valid programming paradigm?
(A) Environmental
(B) Imperative
(C) Declarative
(D) Object-Oriented

57. The property of a type that specifies the set of allowed values it can hold is known as its:
(A) Scope
(B) Lifetime
(C) Domain
(D) Signature

58. What is the purpose of `#include` directive in C/C++?
(A) To include the content of another file into the source file.
(B) To define a macro.
(C) To start a conditional compilation block.
(D) To create an inline function.

59. Consider the following identifiers in the C language:
(a) `_my_variable`
(b) `2ndPlace`
(c) `Total-Sum`
(d) `for`
Which of the above are valid identifiers?
(A) (a) only
(B) (a) and (b) only
(C) (a), (b) and (c) only
(D) (d) only

60. In C++, if a base class has a non-virtual destructor and you delete a derived class object through a base class pointer, what happens?
(A) Only the base class destructor is called.
(B) The behavior is undefined, often leading to a resource leak.
(C) Only the derived class destructor is called.
(D) The compiler issues an error and prevents compilation.

61. What does the term instantiation refer to in object-oriented programming?
(A) Defining a new class.
(B) Inheriting properties from a base class.
(C) The process of creating an instance (an object) of a class.
(D) Hiding the implementation details of a class.

62. Which data structure is typically used to implement the scan-line polygon fill algorithm efficiently?
(A) A stack
(B) A queue
(C) A hash table
(D) An active edge table (AET) sorted by x-coordinates.

63. Which of the following is a key characteristic of a B-Spline curve?
(A) It provides local control, where moving one control point only affects a part of the curve.
(B) The curve passes through all of its control points.
(C) It provides global control, where moving one control point affects the entire curve.
(D) The degree of the curve is always equal to the number of control points.

64. Consider the following statements about C++ references:
(a) A reference is an alias for an already existing variable.
(b) A reference must be initialized at the time of declaration.

Click Here to Buy book on Amazon

class.
(C) A member function that returns a value indicating creation success.
(D) A special member function that is automatically called when an object is created.

28. Which type of polygon filling algorithm is more suitable for complex polygons and generally more efficient because it processes the polygon scan line by scan line?
(A) Boundary-Fill Algorithm
(B) Scan Line Polygon Fill Algorithm
(C) Flood-Fill Algorithm
(D) Edge-Fill Algorithm

29. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : In C++, a function declared as `virtual` in a base class allows for runtime polymorphism.
Reason (R) : When a virtual function is called through a base class pointer pointing to a derived class object, the derived class's version of the function is invoked.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

30. Which HTML tag is used to create a hyperlink?
(A) `<link>`
(B) `<href>`
(C) `<ul>`
(D) `<a>`

31. In 3D object representation, what is a quadric surface?
(A) A surface described by a second-degree equation (e.g., spheres, ellipsoids).
(B) A surface described by a polynomial equation of degree four.
(C) A surface formed by four connected polygons.
(D) A four-sided flat polygon in 3D space.

32. What is the most appropriate description of Sequence Control in programming languages like C?
(A) Managing the order of data bytes in memory.
(B) Determining the order in which statements are executed.
(C) Controlling access to subprograms and functions.
(D) Handling file input and output streams.

33. Which statement about DHTML (Dynamic HTML) is incorrect?
(A) It uses a combination of HTML, CSS, and scripting languages like JavaScript.
(B) It is a standalone language separate from HTML.
(C) It allows a web page to be changed after it has loaded.
(D) It enables dynamic content and user interaction without reloading the page.

34. Match the following C++ concepts in List I with their corresponding keywords or operators in List II.
List I: (Concept)
P. Inheritance
Q. Object Instantiation
R. Exception Handling
S. Polymorphism (Runtime)
List II: (Keyword/Operator)
1. `try`, `catch`
2. `virtual`
3. `:` (colon) in class declaration
4. `new`
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-4, Q-3, R-2, S-1

35. The process of converting a window in world coordinates to a viewport on the display device is called:
(A) Projection
(B) Clipping
(C) Viewing Transformation
(D) Rasterization

36. Consider the following statements regarding operator overloading in C++:
Statement I : In C++, operator overloading allows programmers to redefine the meaning of

Downloaded from www.pdf.tube

CHAPTER 03

Programming Languages and Computer Graphics

1. Which of the following best defines a programming paradigm?
(A) A specific syntax used by a programming language.
(B) A fundamental style or model of computer programming.
(C) The set of tools available in a programming environment.
(D) The process of converting source code to machine code.

2. In the context of programming language concepts, what does binding time refer to?
(A) The moment a name is linked to an object or property.
(B) The total time taken for a program to be compiled.
(C) The duration for which a variable stays in memory.
(D) The time required to execute a specific code function.

3. Which of the following is considered a scalar data type?
(A) An array
(B) A structure
(C) A character
(D) A class

4. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : Static binding, which occurs at compile time, generally leads to more efficient program execution.
Reason (R) : The compiler can generate machine code that directly calls a subprogram or accesses a variable, avoiding the overhead of a runtime lookup.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

5. Which stage of the translation process is primarily responsible for grouping characters into meaningful sequences called lexemes or tokens?
(A) Syntax Analysis
(B) Semantic Analysis
(C) Lexical Analysis
(D) Code Generation

6. Identify the correct statement regarding object-oriented programming (OOP).
(A) Encapsulation allows an object to take many forms.
(B) An object is an instance of a class.
(C) Inheritance describes the bundling of data and methods that operate on the data.
(D) Polymorphism is the mechanism for deriving a new class from an existing one.

7. What is the primary purpose of a virtual function in C++?
(A) To allow static binding of functions.
(B) To enable function overloading.
(C) To create functions that cannot be inherited.
(D) To achieve runtime polymorphism.

8. In computer graphics, which device works by directing an electron beam to the specific parts of the screen where a picture is to be drawn?
(A) Random-scan display
(B) Raster-scan display
(C) Liquid Crystal Display (LCD)
(D) Plasma panel display

45. The reflection of a 2D point (x, y) about the x-axis generates a new point (x', y'). What are the coordinates of (x', y')?
(A) (-x, y)
(B) (-x, -y)
(C) (x, -y)
(D) (y, x)

46. In 3D graphics, which method renders polygons by calculating the intensity at each vertex and interpolating these intensities for the interior pixels?
(A) Phong Shading
(B) Ray Tracing
(C) Z-buffering
(D) Gouraud Shading

47. Identify the correct pair related to C++ I/O streams.
(A) `cerr` : Unbuffered standard error stream
(B) `cin` : Standard output stream
(C) `cout` : Standard input stream
(D) `clog` : Buffered standard input stream

48. A programmer needs to handle a situation where a function might fail to allocate memory. The most appropriate mechanism in C++ to manage this potential failure gracefully is:
(A) Function overloading
(B) Exception handling
(C) Virtual functions
(D) Templates

49. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The parser, or syntax analyzer, uses the grammar of the language to check if the stream of tokens from the lexical analyzer forms a valid program structure.
Reason (R) : The primary output of the syntax analysis stage is an abstract syntax tree (AST).
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

50. Consider the following statements about C++:
(a) A destructor is a member function that is automatically called when an object goes out of scope or is explicitly deleted.
(b) C++ supports multiple inheritance, where a class can inherit from more than one base class.
(c) A template allows writing generic programs where the data type is a parameter.
(d) A virtual function must be defined in the base class and cannot be overridden.
(e) A `const` member function is not allowed to modify any data members of the class.
Which of the above statements are incorrect?
(A) (d) only
(B) (a) and (e) only
(C) (b) and (c) only
(D) (d) and (e) only

51. The boundary-fill algorithm is a method for filling a region in a raster graphics system. It starts from an interior point and requires which two pieces of information?
(A) The fill color and the background color.
(B) The starting point and the fill color.
(C) The fill color and the boundary color.
(D) The starting point and the boundary color.

52. In a raster-scan system, the screen is refreshed by scanning from:
(A) Center to the edges in a spiral.
(B) Random points on the screen.
(C) Right to left, bottom to top.
(D) Top to bottom, one scan line at a time.

53. What is the primary role of a line clipping algorithm like Cohen-Sutherland?
(A) To draw a line using only integer arithmetic.
(B) To smooth the jagged appearance of a line on a raster display.
(C) To determine which part of a line is inside a given rectangular window.
(D) To transform a line from world coordinates to view coordinates.

54. Match the following transformations in List I with their corresponding 2x2 matrix representations in List II (for a 2D system, not homogeneous).
List I: (Transformation)
P. Scaling by (Sx, Sy)

Click Here to Buy book on Amazon

(A) Cannot be used to create objects.
(B) Cannot have any member functions.
(C) Cannot be inherited from.
(D) Contains only private members.

19. Consider the following statements regarding Bezier and B-Spline curves:
(a) A Bezier curve always passes through the first and last control points.
(b) B-Spline curves offer better local control than Bezier curves.
(c) Increasing the degree of a Bezier curve requires changing all basis functions.
(d) The shape of a B-Spline curve is not influenced by the position of all control points globally.
(e) B-Spline curves are less computationally stable than Bezier curves of the same degree.
Which of the above combinations is correct?
(A) (a), (b), and (e) are the only correct items.
(B) (a), (b), (c), and (d) are the only correct items.
(C) (c), (d), and (e) are the only correct items.
(D) (a), (c), and (e) are the only correct items.

20. A programmer wants to ensure that a function `compute()` can be overridden in derived classes and that the decision on which version of `compute()` to call is made at runtime. Which C++ keyword should the programmer use in the base class declaration of `compute()`?
(A) static
(B) const
(C) override
(D) virtual

21. Identify the correct pair of a programming concept and its description.
(A) Instantiation : The process of creating an object of a class.
(B) Encapsulation : Bundling data and methods together.
(C) Inheritance : A mechanism allowing a class to have multiple names.
(D) Polymorphism : The ability of a message to be displayed in more than one form.

22. What is the primary function of a graphics monitor?
(A) To process graphical data from the CPU.
(B) To display information in a visual form.
(C) To store graphical images permanently.
(D) To receive input from the user via a stylus.

23. In the 2D viewing pipeline, what is the role of the Window-to-Viewport coordinate transformation?
(A) To clip objects that are outside the world coordinate limits.
(B) To convert 3D coordinates into 2D screen coordinates.
(C) To transform object coordinates into world coordinates.
(D) To map a selected part of a scene in world coordinates to device coordinates.

24. Arrange the following stages of the programming language translation process in chronological order:
(i) Syntax Analysis
(ii) Code Generation
(iii) Lexical Analysis
(iv) Semantic Analysis
Which of the following represents the correct sequence?
(A) (iii), (i), (iv), (ii)
(B) (i), (iii), (iv), (ii)
(C) (iii), (iv), (i), (ii)
(D) (iv), (iii), (i), (ii)

25. The mid-point circle algorithm determines the pixels for a circle by:
(A) Using trigonometric functions to calculate each point.
(B) Solving the circle equation y = sqrt(r^2 - x^2) for each x.
(C) Using an iterative method based on a decision parameter to select the next pixel.
(D) Transforming a line segment into a circular arc.

26. Which of the following is NOT a valid token in the C language?
(A) A data type
(B) A keyword
(C) An identifier
(D) A string literal

27. In C++, what is a constructor?
(A) A member function used to deallocate memory of an object.
(B) A non-member function used to build a

NODIA

(c) A reference, once initialized, can be made to refer to a different variable later.
(d) A reference uses the same memory address as the original variable.
(e) A `NULL` reference can be created to indicate it points to nothing.
Which of the above combinations is correct?
(A) (a), (b), and (d) only
(B) (c), (d), and (e) only
(C) (a), (c), and (e) only
(D) (b), (c), and (d) only

65. A Java Servlet is a Java class that:
(A) Runs on the client's web browser to create dynamic UI.
(B) Runs on a web server to handle client requests and generate dynamic content.
(C) Is used for standalone graphical applications.
(D) Connects a Java application to a database.

66. What is the Illumination Model in computer graphics used for?
(A) To simulate how light interacts with surfaces to determine their final color and brightness.
(B) To model the shape and geometry of 3D objects.
(C) To clip objects that are outside the viewing volume.
(D) To represent curved surfaces using control points.

67. Consider the following statements regarding polygon filling algorithms:
Statement I : The flood-fill algorithm fills a region by replacing a specified interior color with a new fill color.
Statement II : The boundary-fill algorithm fills a region until it encounters a boundary of a specified color.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

68. The transformation that changes the size of an object is called:
(A) Translation
(B) Rotation
(C) Shearing
(D) Scaling

69. Arrange the following steps of the 3D viewing pipeline in the correct logical order:
(i) Projection Transformation
(ii) Clipping
(iii) Viewport Transformation
(iv) Modeling and Viewing Transformation
Which of the following represents the correct sequence?
(A) (i), (ii), (iii), (iv)
(B) (iv), (i), (ii), (iii)
(C) (iv), (ii), (i), (iii)
(D) (ii), (iv), (iii), (i)

70. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The use of homogeneous coordinates in computer graphics allows for the uniform representation of translation, rotation, and scaling as matrix multiplications.
Reason (R) : Translation, which is an additive operation in Cartesian coordinates, becomes a multiplicative operation in the homogeneous system.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

71. In C++, which of the following correctly describes a destructor?
(A) A member function used to initialize the data members of an object.
(B) A member function that can accept parameters but cannot have a return type.
(C) A global function that is called to create an object.
(D) A member function with the same name as the class, prefixed with a tilde (~).

72. The general form of a projection transform in 3D viewing is responsible for:
(A) Mapping 3D coordinates onto a 2D view plane.
(B) Changing the size of an object in 3D space.
(C) Rotating an object around a specific axis.
(D) Defining the color and material properties

Downloaded from www.pdf.tube

of an object.

73. Which of the following is NOT a polygon rendering method?
(A) Gouraud shading
(B) Phong shading
(C) Cohen-Sutherland shading
(D) Constant intensity shading

74. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : In object-oriented programming, an abstract class cannot be instantiated.
Reason (R) : An abstract class is designed only to be a base class from which other classes may be derived and must contain at least one pure virtual function.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

75. Consider a C program where a function `my_func` is declared as `void my_func(int *p)`. If `x` is an integer variable, which of the following is the correct way to call this function to allow it to modify the value of `x`?
(A) `my_func(x);`
(B) `my_func(*x);`
(C) `my_func(**x);`
(D) `my_func(&x);`

76. In C++, what is the purpose of the `template` keyword?
(A) To define a class that cannot be inherited.
(B) To create functions and classes that can operate with generic types.
(C) To handle runtime errors and exceptions.
(D) To specify the memory layout of an object.

77. Which of the following is an example of a formal transition model used in language theory?
(A) Flowchart
(B) UML Diagram
(C) Finite Automaton
(D) Data Flow Diagram

78. The semantic analysis phase of a compiler is primarily responsible for:
(A) Generating machine code.
(B) Grouping characters into tokens.
(C) Parsing the grammatical structure of the program.
(D) Checking for type mismatches and other meaning-related errors.

79. Match the following web technologies in List I with their primary purposes in List II.
List I: (Web Technology)
P. HTML
Q. XML
R. Java Servlet
S. DHTML
List II: (Primary Purpose)
1. Server-side processing of web requests.
2. Creating interactive and animated web pages.
3. Defining the structure and content of a web page.
4. Storing and transporting data in a platform-independent way.
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-4, Q-3, R-2, S-1

NODIA

80. A programmer encounters a `NullPointerException` while running a Java program. What is the most likely cause?
(A) Attempting to divide a number by zero.
(B) Attempting to use a reference variable that has not been initialized to an object.
(C) Trying to access an array element with an out-of-bounds index.
(D) The program has run out of available memory.

81. Which property of a Bezier curve is defined by the fact that the curve is contained entirely within the convex hull of its control points?
(A) Endpoint interpolation property
(B) Convex hull property
(C) Symmetry property
(D) Local control property

82. In C, what is the key difference between a structure and a union?
(A) Structures allocate memory for all members;

unions allocate only for the largest.
(B) Structures hold different types; unions are restricted to only one shared type.
(C) Structures act as value types; unions represent complex reference types in C.
(D) Structures exclude member functions; unions allow for encapsulated logic.

83. Which of the following best describes multifile programs in the context of C++?
(A) A program that can read from and write to multiple files simultaneously.
(B) A program that includes multiple standard library headers.
(C) A program whose source code is split across several `.cpp` and `.h` files.
(D) A program that can be compiled by different compilers.

84. Consider the following statements about C++ constructors and destructors:
(a) A class can have multiple constructors (overloaded).
(b) A class can have multiple destructors.
(c) A destructor can take arguments.
(d) A constructor is automatically called when an object is created.
Which of the above statements are true?
(A) (a) and (d) only
(B) (a), (b) and (d) only
(C) (b) and (c) only
(D) (a), (c) and (d) only

85. A Java Applet's lifecycle methods include `init()`, `start()`, `stop()`, and `destroy()`. Which method is called only once during the applet's lifetime?
(A) `start()`
(B) `stop()`
(C) `init()`
(D) `paint()`

86. What is the primary advantage of the Mid-point Ellipse Algorithm?
(A) It can draw ellipses of any orientation.
(B) It involves only integer arithmetic, making it computationally efficient.
(C) It uses floating-point calculations for higher precision.
(D) It works by transforming a circle into an ellipse.

87. Which of the following is NOT an input device in the context of computer graphics?
(A) Mouse
(B) Keyboard
(C) Light Pen
(D) Plotter

88. What is the Viewing Coordinate Reference Frame?
(A) A coordinate system defined with the viewer's position as the origin.
(B) The coordinate system in which the display device operates.
(C) The coordinate system defined by the 3D world scene.
(D) The coordinate system used to model individual objects.

89. Consider the following statements about 3D transformations:
(a) A 3D translation requires a 3x3 transformation matrix.
(b) Using 4x4 matrices with homogeneous coordinates allows 3D translation to be represented as a matrix multiplication.
(c) Rotation in 3D can occur around any arbitrary axis, not just the principal axes.
(d) The order of applying multiple transformations (e.g., rotation then translation) does not matter.
(e) 3D shearing distorts the shape of an object, similar to its 2D counterpart.
Which of the above combinations is correct?
(A) (a), (c), and (d) only
(B) (b), (c), and (e) only
(C) (a), (d), and (e) only
(D) (b), (d), and (e) only

90. In C++, when a base class method is declared as `protected`, it can be accessed by:
(A) Members of the base class and its derived classes.
(B) Any function in the program.
(C) Only the members of the base class itself.
(D) Only by `friend` functions of the base class.

91. Which line clipping algorithm uses a 4-bit outcode to classify the endpoints of a line relative to the clipping window?
(A) Liang-Barsky algorithm
(B) Sutherland-Hodgman algorithm

Click Here to Buy book on Amazon

(c) A debugger to find and fix errors.
(d) A specific hardware architecture.
(e) Version control systems like Git.
Which of the above combinations represents the most fundamental components?
(A) (a), (b), and (c) only
(B) (a), (b), (c), and (e) only
(C) (d) and (e) only
(D) (a) and (b) only

110. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The `sizeof` operator in C can be used to determine the size in bytes of a data type or a variable.
Reason (R) : The size of data types like `int` is fixed across all computer architectures and compilers.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

111. The term binding in programming languages refers to the association of an attribute with an entity. When does dynamic binding occur?
(A) At run time
(B) At compile time
(C) At link time
(D) At load time

112. Identify the incorrect statement regarding C structures.
(A) Structure members are accessed using the dot (`.`) operator with a structure variable.
(B) All members of a structure share the same memory location.
(C) A structure can contain members of different data types.
(D) Structures can be nested within other structures.

113. A C++ programmer writes a class to manage a network connection. If the connection is not properly closed, a resource leak will occur. Where is the most appropriate place to put the code that closes the connection to ensure it is always executed when the object is destroyed?
(A) In a public `close()` method that must be called manually.
(B) In the class destructor.
(C) In the class constructor.
(D) In a static member function.

114. Arrange the following 2D viewing pipeline operations in the correct order:
(i) Clip against view window.
(ii) Construct world-coordinate scene using modeling-coordinate transformations.
(iii) Transform from world coordinates to viewing coordinates.
(iv) Map viewing coordinates to normalized viewport coordinates.
Which of the following represents the correct sequence?
(A) (ii), (iii), (i), (iv)
(B) (ii), (i), (iii), (iv)
(C) (iii), (ii), (iv), (i)
(D) (i), (ii), (iii), (iv)

115. What is the transformation matrix for a 2D rotation by an angle θ about the origin?
(A) `[[cosθ, sinθ], [-sinθ, cosθ]]`
(B) `[[sinθ, cosθ], [cosθ, -sinθ]]`
(C) `[[cosθ, -sinθ], [sinθ, cosθ]]`
(D) `[[1, 0], [0, 1]]`

116. Consider the following statements regarding spline curves:
Statement I : A Bezier curve of degree n is defined by n control points.
Statement II : A primary advantage of B-spline curves over Bezier curves is local control.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

117. Which of the following is a key component of the Phong illumination model?
(A) Z-buffer
(B) Scan line
(C) Specular reflection
(D) Outcode

118. Consider the following statements about C++:
(a) `this` is a pointer accessible within a member

(d) Operator overloading
(e) Namespaces
Which of the following features is most appropriate?
(A) (a)
(B) (b)
(C) (c)
(D) (d)

101. What is the fundamental difference between a raster-scan and a random-scan display?
(A) Raster draws one line at a time; random draws point by point.
(B) Raster requires a shadow mask; random scan does not need one.
(C) Raster targets the whole screen; random draws only where needed.
(D) Raster is for monochrome displays; random is for color monitors.

102. The process of mapping a part of a world-coordinate scene to device coordinates is known as the:
(A) Modeling transformation
(B) Viewing transformation
(C) Projection transformation
(D) Normalization transformation

103. In C++, what is the purpose of overloading a function?
(A) To allow a function in a derived class to have the same name as a function in the base class.
(B) To prevent a function from being called.
(C) To create a function that can operate on different data types using a single definition.
(D) To allow multiple functions with the same name but different parameter lists to exist.

104. Consider the following statements about C++ file streams:
(a) `ofstream` is a stream class to write on files.
(b) `ifstream` is a stream class to read from files.
(c) `fstream` is a stream class that supports both read and write operations.
(d) The `<iostream>` header must be included for file operations.
Which of the above statements are true?
(A) (a) and (b) only
(B) (a), (b) and (c) only
(C) (c) and (d) only
(D) (a), (b), (c) and (d)

105. In a scripting language like JavaScript, what is the role of the Document Object Model (DOM)?
(A) Server-side interface for connecting and accessing databases.
(B) Security framework to block scripts from running in browsers.
(C) Interface to access and modify document structure and content.
(D) Compiler component for turning scripts into machine-level code.

106. In 3D graphics, a projection that preserves parallel lines as parallel in the projected image is known as:
(A) Orthographic Projection
(B) Perspective Projection
(C) Stereoscopic Projection
(D) Central Projection

107. Which of the following best describes spline representation in computer graphics?
(A) A method for representing 3D objects using a mesh of polygons.
(B) A technique for filling the interior of a polygon.
(C) An algorithm for clipping lines against a rectangular window.
(D) A mathematical representation of a curve or surface defined by a set of control points.

NODIA

108. The Phong illumination model differs from the Gouraud model primarily because:
(A) Phong model calculates lighting at each vertex, while Gouraud calculates it for every pixel.
(B) Phong model interpolates surface normals across a polygon and calculates lighting for each pixel, producing better specular highlights.
(C) Gouraud model includes ambient, diffuse, and specular components, while Phong only includes diffuse.
(D) Phong model is computationally less expensive than the Gouraud model.

109. Consider the following components of a programming environment:
(a) A text editor for writing code.
(b) A compiler or interpreter to translate code.

R. Boundary-Fill Algorithm
S. Gouraud Shading

List II: (Primary Function)

1. Polygon filling by starting from an interior point.
2. Line clipping against a rectangular window.
3. Efficient line drawing using integer arithmetic.
4. Smooth shading by interpolating color intensities.

(A) P-1, Q-4, R-2, S-3
(B) P-3, Q-2, R-1, S-4
(C) P-3, Q-1, R-2, S-4
(D) P-2, Q-3, R-1, S-4

145. What happens during the syntax analysis or parsing stage of compilation?
(A) The meaning of the program is checked for logical errors.
(B) The source code is checked to ensure it conforms to the grammatical rules of the language.
(C) The source code is broken down into a stream of tokens.
(D) Machine-independent intermediate code is generated.

146. Which of the following is an example of a scalar data type?
(A) `float temperature`
(B) `struct customer`
(C) `int student_marks[50]`
(D) `union data`

147. In C, what will `int *p;` declare?
(A) An integer variable `p`.
(B) A function `p` that returns a pointer to an integer.
(C) A pointer `p` to an integer.
(D) A pointer `p` to a function.

148. Which OOP principle is demonstrated by having a `Shape` base class with a `draw()` method, and derived classes like `Circle` and `Square` that implement `draw()` in their own way?
(A) Encapsulation
(B) Inheritance
(C) Abstraction
(D) Polymorphism

149. Consider the following keywords in C++:
(a) `public`
(b) `private`
(c) `protected`
(d) `virtual`
Which of the above is NOT an access specifier for class members?
(A) (d) only
(B) (c) and (d) only
(C) (b) only
(D) (a), (b) and (c)

150. Consider the following HTML tags used in form creation:
(a) `<input>`
(b) `<form>`
(c) `<button>`
(d) `<table>`
(e) `<label>`
Which of the above provides the container for the entire form?
(A) (a)
(B) (b)
(C) (c)
(D) (d)

151. Which of the following devices is a classic example of a random-scan system?
(A) A modern LCD monitor
(B) A CRT television
(C) A laser printer
(D) A vector-refresh display

152. In 2D transformations, the matrix for reflecting a point across the line y = x is:
(A) `[[0, 1], [1, 0]]`
(B) `[[1, 0], [0, -1]]`
(C) `[[-1, 0], [0, 1]]`
(D) `[[0, -1], [-1, 0]]`

153. What is the Viewing Coordinate Reference Frame (VCRF) defined by?
(A) A viewport on the screen and a window in the world.
(B) The boundaries of the 3D world.
(C) A view point (eye), a look-at point, and a view-up vector.
(D) The resolution of the display device.

154. Consider the following statements regarding 3D projections:
(a) Perspective projections are more realistic as

Downloaded from www.pdf.tube

function that points to the object for which the function was called.
(b) A static member function can directly access non-static data members of the class.
(c) Overloading the `[]` operator allows objects to be indexed like arrays.
(d) A friend function of a class can access the private and protected members of that class.
(e) `new` and `delete` are functions, not operators.
Which of the above statements are incorrect?
(A) (b) and (e) only
(B) (a) and (c) only
(C) (d) and (e) only
(D) (b) and (d) only

119. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : Encapsulation is a fundamental principle of OOP that involves bundling the data (attributes) and methods (functions) that operate on the data into a single unit or class.
Reason (R) : The primary goal of encapsulation is to facilitate code reuse through inheritance.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

120. In the C `main` function `int main(int argc, char *argv[])`, what does `argc` represent?
(A) The number of command line arguments plus one (for the program name).
(B) The name of the program.
(C) A pointer to the array of command line arguments.
(D) The return code of the program.

121. Which programming paradigm is characterized by its use of facts, rules, and queries?
(A) Object-Oriented Programming
(B) Functional Programming
(C) Logic Programming
(D) Procedural Programming

122. Which of the following is an example of a composite data type?
(A) `struct`
(B) `char`
(C) `float`
(D) `enum`

123. A function in C is defined as `int sum(int a, int b) { return a+b; }`. This mechanism of function definition and calling falls under the category of:
(A) Sequence Control
(B) File Handling
(C) Preprocessing
(D) Subprogram Control

NODIA

124. Match the following Object-Oriented Programming (OOP) concepts in List I with their descriptions in List II.
List I: (OOP Concept)
P. Class
Q. Object
R. Method
S. Attribute
List II: (Description)
1. A function or procedure associated with a class.
2. A data member or property of a class.
3. A blueprint or template for creating objects.
4. A specific instance of a class.
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-4, R-2, S-1
(D) P-1, Q-2, R-3, S-4

125. Which C++ feature allows a derived class to provide a specific implementation of a method that is already provided by its base class?
(A) Function Overloading
(B) Operator Overloading
(C) Function Overriding
(D) Function Templating

126. What is the primary difference between HTML and XML?
(A) HTML is case-sensitive, while XML is not.
(B) HTML is used for data transport, while XML is for displaying data.
(C) HTML requires a closing tag for every opening tag, while XML does not.
(D) HTML tags are predefined, while XML allows you to define your own tags.

(C) Cyrus-Beck algorithm
(D) Cohen-Sutherland algorithm

92. Which of the following statements about polygon surfaces is incorrect?
(A) They are commonly used to represent the surfaces of 3D objects.
(B) A polygon surface must always be convex.
(C) They are defined by a set of vertex coordinates and an attribute list.
(D) Rendering of polygon surfaces can involve filling and shading.

93. What is a key characteristic of a B-spline surface?
(A) It is defined by a single patch of control points.
(B) It must interpolate all of its control points.
(C) It provides local control over the surface shape.
(D) It can only represent flat, planar surfaces.

94. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The `virtual` keyword is essential for achieving runtime polymorphism in C++.
Reason (R) : When a derived class object is pointed to by a base class pointer, the `virtual` keyword ensures that the correct overridden function in the derived class is called.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

95. In the context of programming languages, what is a token?
(A) The smallest individual unit of a program, such as a keyword, identifier, or operator.
(B) A variable name.
(C) A comment in the source code.
(D) A complete executable statement.

96. Which C preprocessor directive is used for conditional compilation?
(A) `#include`
(B) `#ifdef`
(C) `#define`
(D) `#pragma`

97. What is the main difference between call-by-value and call-by-reference parameter passing?
(A) Call-by-value passes the address of the argument, while call-by-reference passes the value.
(B) Call-by-value passes a copy of the argument's value, while call-by-reference passes the argument itself, allowing modification.
(C) Call-by-value is used for integers, while call-by-reference is used for pointers.
(D) Call-by-value is more memory efficient than call-by-reference.

98. Which of the following is NOT a type of inheritance supported by C++?
(A) Delegated Inheritance
(B) Single Inheritance
(C) Multiple Inheritance
(D) Hierarchical Inheritance

99. Match the following C++ features in List I with their primary benefits in List II.
List I: (C++ Feature)
P. Encapsulation
Q. Inheritance
R. Polymorphism
S. Template
List II: (Primary Benefit)
1. Code Reusability and extending functionality.
2. Writing type-independent (generic) code.
3. Flexibility and dynamic behavior via a common interface.
4. Information hiding and data protection.
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-1, R-3, S-2
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-1, R-4, S-3

100. A programmer wants to write a C++ function that can compute the average of an array of integers, an array of floats, or an array of doubles without writing three separate functions. Consider the following C++ features:
(a) Function overloading
(b) Virtual functions
(c) Templates

Click Here to Buy book on Amazon

136. Which of the following is a key property of an object in the context of Elementary Data Types?
(A) Its type, which determines the operations that can be performed on it.
(B) Its programming paradigm.
(C) The name of the file it is defined in.
(D) The programmer who wrote the code for it.

137. In C programming, `char *str = Hello;` declares `str` as a:
(A) Character array of size 6.
(B) Single character variable.
(C) Pointer to a constant string literal.
(D) Structure containing a string.

138. Which of the following best defines polymorphism in OOP?
(A) The ability to process objects of different types through a single, uniform interface.
(B) The ability of a class to be derived from more than one base class.
(C) The process of hiding implementation details from the user.
(D) The creation of an object from a class definition.

139. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The Sutherland-Hodgman polygon clipping algorithm works by processing a polygon's vertices against each edge of a convex clip window sequentially.
Reason (R) : The algorithm is guaranteed to produce a single polygon as output, even if the clipping action should result in multiple separate polygons.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

140. Consider the following statements about virtual functions in C++:
(a) They are declared using the `virtual` keyword in the base class.
(b) They enable static (compile-time) polymorphism.
(c) They are used to implement function overloading.
(d) When called through a base class pointer, the version corresponding to the actual object type is invoked.
Which of the above statements are correct?
(A) (a) and (b) only
(B) (a) and (d) only
(C) (b) and (c) only
(D) (c) and (d) only

141. A `try-catch` block in C++ is used for:
(A) Defining a template class.
(B) Overloading operators.
(C) Handling exceptions and preventing program termination from runtime errors.
(D) Implementing virtual functions.

NODIA

142. The `Window to View-Port Coordinate Transformation` involves which two main operations?
(A) Rotation and Shearing
(B) Reflection and Projection
(C) Clipping and Filling
(D) Scaling and Translation

143. Consider the following statements regarding 3D object representations:
(a) Polygon surfaces approximate object surfaces with a mesh of planar facets.
(b) Quadric surfaces like spheres and cylinders are represented by second-degree equations.
(c) A Bezier surface is defined by a rectangular mesh of control points.
(d) Spline representations are generally unsuitable for designing smooth, curved shapes.
(e) B-Spline surfaces offer less local control compared to Bezier surfaces.
Which of the above combinations is correct?
(A) (a), (b), and (c) only
(B) (c), (d), and (e) only
(C) (a), (d), and (e) only
(D) (b), (c), and (d) only

144. Match the following graphics algorithms in List I with their primary functions in List II.
List I: (Graphics Algorithm)
P. Bresenham's Line Algorithm
Q. Cohen-Sutherland Algorithm

127. Which of the following is NOT a graphics primitive?
(A) Color
(B) Point
(C) Line
(D) Polygon

128. What is the main purpose of the Liang-Barsky line clipping algorithm?
(A) To simplify polygon meshes for faster rendering.
(B) To determine the visible portions of a line that are inside a convex clipping window using parametric equations.
(C) To fill the interior of a closed shape with a solid color.
(D) To apply shading to a 3D model.

129. Consider the following statements about C++ inheritance:
(a) Private members of a base class are directly accessible in the derived class.
(b) `public` inheritance establishes an is-a relationship.
(c) A derived class inherits constructors and destructors from its base class.
(d) `protected` members of a base class are accessible in the derived class but not by outside code.
Which of the above statements are correct?
(A) (a) and (c) only
(B) (b) and (d) only
(C) (a), (b) and (d) only
(D) (c) and (d) only

130. A programmer is designing a graphical user interface and wants to ensure that a button can be clicked to perform an action. This requires handling user-initiated actions. This concept is broadly known as:
(A) Event Handling
(B) Exception Handling
(C) Stream Handling
(D) Multifile Programming

131. Which of the following is a defining characteristic of a random-scan (vector) display?
(A) It has a high refresh rate suitable for realistic images.
(B) The entire screen content is stored in a frame buffer.
(C) It draws lines and characters directly, leading to smooth lines.
(D) It is based on liquid crystal technology.

132. The transformation that moves an object from one position to another in a straight line is called:
(A) Scaling
(B) Rotation
(C) Reflection
(D) Translation

133. What is the purpose of the viewing pipeline in computer graphics?
(A) To connect input devices like a mouse and keyboard to the graphics system.
(B) To manage memory allocation for graphical objects.
(C) To define the sequence of operations for creating a 2D image from a 3D scene description.
(D) To store the final image for display.

134. Consider the following statements regarding the functionality of polygon clipping algorithms:
(a) The Sutherland-Hodgman algorithm clips polygons against a convex clipping window.
(b) It processes the polygon one edge at a time against the entire clip window.
(c) It can sometimes produce multiple disjoint polygons from a single clipped polygon.
(d) The Weiler-Atherton algorithm can correctly clip both convex and concave polygons.
(e) Polygon clipping is a simpler process than line clipping.
Which of the above combinations is correct?
(A) (a) and (d) only
(B) (b), (c), and (e) only
(C) (a), (c), and (d) only
(D) (a), (b), and (e) only

135. What is a programming environment?
(A) The hardware on which a program runs.
(B) The operating system under which a program executes.
(C) The paradigm a programming language follows.
(D) A collection of tools and software to aid in the development of programs.

Downloaded from www.pdf.tube

objects farther away appear smaller.
(b) Orthographic projections are often used in engineering and architectural drawings.
(c) In a perspective projection, parallel lines in 3D space may converge to a vanishing point.
(d) Oblique projection is a type of perspective projection.
(e) The center of projection is at infinity for an orthographic projection.

Which of the above combinations is correct?
(A) (a), (b), (c), and (e) only
(B) (c), (d), and (e) only
(C) (a), (d), and (e) only
(D) (a), (b), and (d) only

155. What is the final stage in the language translation process before the output file is generated?
(A) Lexical Analysis
(B) Syntax Analysis
(C) Semantic Analysis
(D) Code Generation

156. Consider the following statements regarding data types:
Statement I : The properties of a data type include the set of values it can represent and the set of operations that can be performed on those values.
Statement II : A composite data type is one that is built up from scalar or other composite types.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

157. What is the output of the following C code snippet?
`int main() { char str[] = PROGRAM; printf(%d, %d, sizeof(str), strlen(str)); return 0; }`
(A) 7, 7
(B) 7, 8
(C) 8, 7
(D) 8, 8

158. A class `Vehicle` is defined. Then, classes `Car`, `Truck`, and `Bus` are all derived from `Vehicle`. This is an example of which type of inheritance?
(A) Multiple Inheritance
(B) Multilevel Inheritance
(C) Hybrid Inheritance
(D) Hierarchical Inheritance

159. Arrange the following C++ features in a logical order of dependency, where one concept is often a prerequisite for the next:
(i) Class
(ii) Object
(iii) Inheritance
(iv) Virtual Function

Which of the following represents the correct sequence?
(A) (i), (iii), (ii), (iv)
(B) (i), (ii), (iii), (iv)
(C) (ii), (i), (iv), (iii)
(D) (iv), (iii), (ii), (i)

160. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The `try...catch` block in C++ provides a mechanism for handling runtime errors in a controlled and structured manner.
Reason (R) : When an error occurs in a `try` block, an exception is thrown, and program control transfers to the first matching `catch` block.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

161. What does DHTML stand for?
(A) Document Hypertext Markup Language
(B) Dynamic Hypertext Markup Language
(C) Distributed Hypertext Markup Language
(D) Descriptive Hypertext Markup Language

162. The Mid-point circle algorithm is based on testing the sign of a decision parameter, which is derived from:
(A) The implicit form of the circle equation, $f(x, y) = x^2 + y^2 - r^2 = 0$.
(B) The Taylor series expansion of the circle equation.
(C) The parametric form of the circle equation, $x = r\cos(\theta)$, $y = r\sin(\theta)$.
(D) The line equation $y = mx + c$.

163. Consider the following statements regarding the characteristics of Bezier curves:
(a) The degree of a Bezier curve is one less than the number of control points.
(b) The curve is always tangent to the first and last segments of the control polygon.
(c) A cubic Bezier curve is defined by 3 control points.
(d) Moving a single control point affects the shape of the entire curve (global control).
(e) The curve always passes through all of its control points.
Which of the above combinations is correct?
(A) (a), (b), and (d) only
(B) (c), (d), and (e) only
(C) (a), (c), and (e) only
(D) (b), (c), and (d) only

164. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : A Java applet runs in a sandbox security model within the client's web browser.
Reason (R) : The sandbox restricts the applet's access to the local file system and network resources to prevent malicious activity.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

165. The `Viewing Coordinate Reference Frame` is also known as:
(A) World Coordinate System
(B) Device Coordinate System
(C) Eye Coordinate System
(D) Object Coordinate System

166. Which polygon rendering method interpolates the normal vector at each pixel and applies the illumination model, resulting in high-quality specular highlights?
(A) Constant Shading
(B) Phong Shading
(C) Gouraud Shading
(D) Scan-line Shading

167. Identify the incorrect pair of a C token and its category.
(A) `printf` : Identifier (library function name)
(B) `+` : Operator
(C) `12.5` : Constant
(D) `auto` : Preprocessor

168. Arrange the following memory areas in C programming based on their typical lifetime, from shortest to longest:
(i) Global variables
(ii) Static variables within a function
(iii) Local variables within a function (automatic)
Which of the following represents the correct order?
(A) (iii), (ii), (i)
(B) (i), (ii), (iii)
(C) (iii), (i), (ii)
(D) (ii), (iii), (i)

169. Match the following C++ terms in List I with their definitions in List II.
List I: (C++ Term)
P. Constructor
Q. Destructor
R. `this` pointer
S. `friend` function
List II: (Definition)
1. A function that can access private members of a class, though it is not a member of the class.
2. A pointer that holds the address of the current object.
3. A member function automatically called to clean up before an object is destroyed.
4. A member function automatically called to initialize an object when it is created.
(A) P-4, Q-3, R-2, S-1
(B) P-3, Q-4, R-2, S-1
(C) P-4, Q-3, R-1, S-2
(D) P-1, Q-2, R-3, S-4

170. Consider the following statements about C++ templates:
(a) Templates are a feature for implementing runtime polymorphism.
(b) The `template` keyword is used to declare a generic class or function.
(c) When a template is used, the compiler generates the appropriate code for the specified data type (instantiation).

Click Here to Buy book on Amazon

(d) Templates reduce code duplication.
Which of the above statements are incorrect?
(A) (a) only
(B) (b) and (c) only
(C) (a) and (d) only
(D) (c) and (d) only

171. What is the fundamental building block of a raster-scan display's image?
(A) A vector
(B) A scan line
(C) A character
(D) A pixel

172. A 2D shear transformation parallel to the x-axis is represented by the matrix `[[1, shx], [0, 1]]`. What does the parameter `shx` control?
(A) The amount of vertical displacement.
(B) The degree to which x-coordinates are shifted based on the y-coordinate.
(C) The scaling factor along the x-axis.
(D) The angle of rotation around the z-axis.

173. What is the primary purpose of clipping in computer graphics?
(A) To increase the resolution of an image.
(B) To apply color and shading to objects.
(C) To identify and discard parts of a picture that are outside the viewing window or volume.
(D) To combine multiple transformation matrices into one.

174. Consider the following characteristics of programming language syntax:
(a) Readability and writability.
(b) Lack of ambiguity.
(c) Orthogonality (a small set of primitive constructs can be combined in a small number of ways).
(d) Close resemblance to machine code for efficiency.
Which of the above statements are correct?
(A) (a), (b), and (c) only
(B) (a) and (d) only
(C) (c) and (d) only
(D) (a), (b), (c), and (d)

175. Consider the following statements regarding elementary data types:
(a) A scalar type is also known as a simple or primitive type.
(b) A composite type is constructed from one or more other types.
(c) A `struct` in C is an example of a scalar type.
(d) The property of a type includes its name, value range, and the operations defined on it.
(e) The type of an object determines how its value is stored in memory.
Which of the above combinations is correct?
(A) (a), (b), (c), (d), (e)
(B) (a), (b), (d), (e) only
(C) (c), (d), (e) only
(D) (a), (c), (d) only

176. What is the purpose of the `NULL` pointer in C?
(A) To point to the beginning of the program memory.
(B) To indicate that the pointer does not point to any valid memory location.
(C) To create a pointer that cannot be modified.
(D) To store the address of the `main` function.

177. If class `D` is derived from class `B`, and class `B` is derived from class `A`, this is an example of:
(A) Multiple inheritance
(B) Multilevel inheritance
(C) Hierarchical inheritance
(D) Hybrid inheritance

178. In C++, what is the `::` operator used for?
(A) To specify a class's scope for a member or to access global variables (Scope Resolution Operator).
(B) To access members of an object through a pointer.
(C) To declare a reference variable.
(D) To perform pointer arithmetic.

179. Consider the following Object-Oriented Programming (OOP) concepts:
(a) Inheritance
(b) Polymorphism
(c) Encapsulation
(d) Virtual functions
(e) Constructor overloading
Which of the above combination is most directly relevant to achieving runtime dynamic behavior?
(A) (a), (b), and (d)
(B) (c) and (e)
(C) (a) and (c)

Downloaded from www.pdf.tube

(D) Code Generation

230. Match the following programming concepts in List I with their descriptions in List II.
List I: (Programming Concept)
P. Type Checking
Q. Binding Time
R. Paradigm
S. Syntax
List II: (Description)
1. The grammatical rules that govern how symbols may be combined to create a well-formed program.
2. The fundamental style or model of programming.
3. The process of ensuring that an operator or function is applied to operands of the correct type.
4. The point at which an entity is associated with an attribute, such as a variable to its type.
(A) P-3, Q-4, R-2, S-1
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-1, R-2, S-4
(D) P-1, Q-2, R-3, S-4

231. In C, if `p` is a pointer to a structure, which operator is used to access its members?
(A) `.` (dot)
(B) `*` (asterisk)
(C) `&` (ampersand)
(D) `->` (arrow)

232. Which OOP concept allows a class to hide its internal state and require all interaction to be done through an object's methods?
(A) Encapsulation
(B) Inheritance
(C) Polymorphism
(D) Abstraction

233. What is the function of the `new` operator in C++?
(A) It allocates memory from the stack for an object.
(B) It only calls the constructor of an object without allocating memory.
(C) It allocates memory from the heap for an object and calls its constructor.
(D) It deallocates memory for an object.

234. In web programming, which technology runs on the client-side to make web pages interactive?
(A) Scripting (e.g., JavaScript)
(B) Servlets
(C) JavaServer Pages (JSP)
(D) XML

235. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : In a random-scan display, the electron beam is directed only to the parts of the screen where a picture is to be drawn.
Reason (R) : Ideal for line drawing but unsuitable for realistic shading.
(A) Produces smooth lines but lacks efficiency for complex scenes.
(B) Ideal for line drawing but unsuitable for realistic shading.
(C) Requires a high-density frame buffer to store pixels properly.
(D) Relies on liquid crystal technology for displaying vector data.

NODIA

236. A composite transformation matrix is formed by:
(A) Adding the individual transformation matrices.
(B) Multiplying the individual transformation matrices.
(C) Averaging the elements of the individual matrices.
(D) Finding the inverse of each individual matrix.

237. Which of the following is NOT a property of B-spline curves?
(A) The curve generally follows the shape of the control polygon.
(B) They exhibit local control.
(C) The curve must pass through the first and last control points.
(D) The degree of the curve is independent of the number of control points.

238. Consider the following statements regarding constant pointers in C++:
Statement I : `const int *p;` declares a pointer `p` to a constant integer. The integer's value

(C) Quality and extensibility of the provided standard library.
(D) Small set of primitives combined to build complex structures.

188. Consider the following approaches to pass a large struct to a function in C:
(a) Pass the structure by value.
(b) Pass a pointer to the structure.
(c) Declare the structure as global.
(d) Convert the structure to an array before passing.
(e) Pass each member of the structure as a separate argument.
Which of the above is the most efficient and conventional method?
(A) (a)
(B) (b)
(C) (c)
(D) (e)

189. Match the following C++ operators in List I with their typical use cases in List II.
List I: (C++ Operator)
P. `->`
Q. `.`
R. `::`
S. `&`
List II: (Typical Use Case)
1. Accessing a member of an object variable.
2. Accessing a member of an object through a pointer.
3. Getting the memory address of a variable.
4. Scope resolution (e.g., accessing a static member).
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-4, R-1, S-3
(D) P-4, Q-1, R-2, S-3

190. Consider the following statements about C++ exception handling:
(a) The `try` block contains code that may throw an exception.
(b) A single `try` block can be associated with multiple `catch` blocks.
(c) The `throw` keyword is used to raise an exception.
(d) Once an exception is caught, the program execution resumes at the statement immediately following the `throw` statement.
Which of the above statements are incorrect?
(A) (d) only
(B) (a) and (c) only
(C) (b) and (d) only
(D) (c) only

191. A video display device that stores the picture information as a charge distribution on a grid of insulated capsules is a:
(A) Plasma Panel
(B) Cathode Ray Tube (CRT)
(C) Light Emitting Diode (LED) display
(D) Vector Refresh Display

192. Which algorithm is more general for line clipping against a convex polygon clipping window?
(A) Cohen-Sutherland
(B) Cyrus-Beck
(C) Liang-Barsky
(D) Sutherland-Hodgman

193. What is the primary purpose of a `clipping` operation in the 3D viewing pipeline?
(A) To convert 3D coordinates to 2D screen coordinates.
(B) To remove objects or parts of objects that are outside the view frustum.
(C) To apply lighting and shading models to the surfaces.
(D) To assemble the final pixel colors in the frame buffer.

194. Consider the following variable declarations in C:
(a) `int x = 10;`
(b) `const float PI = 3.14;`
(c) `char name[] = test;`
(d) `struct { int a; float b; } record;`
Which of the above are valid variable declarations?
(A) (a) only
(B) (a) and (c) only
(C) (a), (c) and (d) only
(D) (a), (b), (c) and (d)

195. What is the role of a formal transition model in language design?
(A) To provide a graphical representation of program flow.

Click Here to Buy book on Amazon

247. What is the fundamental difference between a perspective projection and a parallel (orthographic) projection?
(A) Parallel projection preserves straight lines, while perspective projection does not.
(B) Parallel projection is used for 2D graphics, while perspective projection is for 3D.
(C) Perspective projection requires homogeneous coordinates, while parallel projection does not.
(D) Perspective projection makes objects appear smaller as their distance from the viewer increases, while parallel projection does not.

248. Match the following C functions in List I with the header files in List II where they are typically declared.
List I: (C Function)
P. `printf()`
Q. `malloc()`
R. `strcpy()`
S. `pow()`
List II: (Header File)
1. `string.h`
2. `math.h`
3. `stdio.h`
4. `stdlib.h`
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-1, Q-4, R-3, S-2

249. Which of the following best describes an abstract class in C++?
(A) A class with no data members.
(B) A class that can only be inherited privately.
(C) A class that has at least one pure virtual function, and cannot be instantiated.
(D) A class whose objects are created on the stack instead of the heap.

250. Consider the following statements regarding the C++ new and delete operators:
(a) `new` allocates memory from the free store (heap).
(b) `new` automatically calls the appropriate constructor for the object being created.
(c) `delete` deallocates memory allocated by `new`.
(d) `delete` automatically calls the appropriate destructor.
(e) `new` and `delete` can be overloaded for custom memory management.
Which of the above combination of statements is correct?
(A) Statements (a), (c), and (e) are correct.
(B) Statements (b), (d), and (e) are correct.
(C) Statements (a), (b), and (d) are correct.
(D) Statements (a), (b), (c), (d), and (e) are correct.

Direction : Please read the following source carefully.

The design and translation of programming languages are foundational to computer science, encompassing a range of concepts, paradigms, and practical models. Programming language concepts refer to the fundamental ideas that shape a language, such as its data typing model, control structures, and memory management approach. These concepts are organized into paradigms, which are distinct schools of thought on how programming should be approached; examples include procedural, object-oriented, and functional programming. A programming environment provides a suite of tools to support the development process, often built around a specific paradigm. A crucial theoretical concept is the virtual computer, an abstract machine for which a language is designed, simplifying the translation process by providing a consistent target. The concept of binding times is central to understanding how language attributes, such as the association of a variable with its type or memory location, are fixed. This can occur at various points, from language design time to run time. The entire process is governed by the language's syntax, which is formally specified. Translation from a high-level language to machine code occurs in distinct stages, typically including lexical analysis, syntax analysis, semantic analysis, intermediate code generation, optimization, and final code generation. These stages are often guided by formal transition models, such as finite automata for lexical analysis and pushdown automata for syntax analysis, ensuring a rigorous and predictable translation process.

Based on the above source, attempt the following five questions.

conceptual order?
(A) (i) and (iii) are equally parallel, followed by (ii)
(B) (ii), (iii), (i)
(C) (iii), (i), (ii)
(D) (i), (ii), (iii)

205. Match the following terms in List I with the most appropriate descriptions in List II.
List I: (Term)
P. Abstract Class
Q. Virtual Function
R. Template
S. Constructor
List II: (Description)
1. A function whose implementation is provided at runtime.
2. A blueprint for creating generic classes or functions.
3. A class that cannot be instantiated and is meant for inheritance.
4. A special method for creating and initializing an object.
(A) P-1, Q-3, R-2, S-4
(B) P-3, Q-1, R-2, S-4
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-2, S-3

206. Consider the following statements regarding binding in C++:
Statement I : Late binding is typically implemented using virtual function tables (vtables).
Statement II : Static binding is resolved at compile time and is generally more efficient than late binding.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

207. Which C library function is used to open a file?
(A) `fopen()`
(B) `fget()`
(C) `openfile()`
(D) `fileno()`

208. In C++, if you have a pointer `p` to an object, how do you access a member `data` of that object?
(A) `p.data`
(B) `p->data`
(C) `p::data`
(D) `p&data`

209. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : XML is a markup language that separates data from presentation.
Reason (R) : XML files contain data and the instructions on how to display that data, making it a self-contained presentational format.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

210. Which of the following best describes a scripting language?
(A) A language that is typically interpreted within another program or environment to automate tasks.
(B) A language that must be compiled to machine code before execution.
(C) A language designed for writing operating systems.
(D) A low-level language that directly manipulates hardware.

211. What is a Viewing Function in a graphics API like OpenGL?
(A) A function to set the color of an object.
(B) A function to draw a geometric primitive like a line or polygon.
(C) A function to define the position and orientation of the camera or viewer.
(D) A function to close the graphics display window.

212. Which of the following is an example of a quadric surface?
(A) A cube
(B) A torus
(C) A Bezier patch
(D) An ellipsoid

Downloaded from www.pdf.tube

(B) To serve as a programming environment for developers.
(C) To provide a mathematical or abstract model for the behavior of a computational system, often used in compiler design.
(D) To specify the hardware requirements for a programming language.

196. The binding of a method call to a specific method implementation at runtime rather than at compile time is called:
(A) Static Binding
(B) Early Binding
(C) Late Binding
(D) Constant Binding

197. Which of the following correctly describes a C preprocessor?
(A) Compiler component that performs static type-checking routines.
(B) System tool that links multiple object files into an executable.
(C) Program that handles directives before the actual compilation.
(D) Software debugger used to step through program code execution.

198. In C++, if a class `C` is a `friend` of class `P`, what does this imply?
(A) Class `P` can access the private and protected members of class `C`.
(B) Class `C` can access the private and protected members of class `P`.
(C) Objects of class `C` and `P` cannot be created.
(D) Class `C` is a derived class of class `P`.

199. Consider the following HTML and scripting concepts:
(a) HTML is responsible for the structure and content of a web page.
(b) CSS (Cascading Style Sheets) is used for the presentation and styling.
(c) JavaScript is a client-side scripting language used to make web pages interactive.
(d) XML is a language primarily used for displaying complex graphical data.
(e) DHTML is a combination of HTML, CSS, and scripting to create dynamic pages.
Which of the above combinations is correct?
(A) (a), (b), (c), and (e) only
(B) (c), (d), and (e) only
(C) (a), (d), and (e) only
(D) (a), (b), (d), and (e) only

200. What is a Bezier surface?
(A) A parametric surface that is an extension of the Bezier curve, defined by a mesh of control points.
(B) A 3D surface generated by sweeping a 2D curve along a path.
(C) A surface represented by an array of height values on a grid.
(D) A surface defined by a second-degree polynomial equation.

201. The composition of two 2D rotations with angles θ1 and θ2 is equivalent to a single rotation with an angle of:
(A) θ1 * θ2
(B) θ1 / θ2
(C) θ1 + θ2
(D) max(θ1, θ2)

202. What is the purpose of the Z-buffer algorithm in 3D graphics?
(A) To determine which surfaces are visible to the viewer at each pixel location.
(B) To clip polygons against the view frustum.
(C) To apply texture maps to surfaces.
(D) To perform smooth shading.

NODIA

203. Consider the following design patterns:
(a) Factory
(b) Singleton
(c) Observer
(d) Decorator
Which of the above design patterns is being implemented to ensure only one instance of a class exists?
(A) (a)
(B) (b)
(C) (c)
(D) (d)

204. Arrange the following general projection types from most parallel (less perspective) to most perspective:
(i) Orthographic Projection
(ii) Perspective Projection
(iii) Axonometric Projection (a type of parallel projection)
Which of the following represents the correct

(A) A derived class can access the public members of its base class.
(B) A derived class can access the protected members of its base class.
(C) A derived class can access the private members of its base class.
(D) C++ supports multiple inheritance.

222. Which of the following is NOT an advantage of using XML?
(A) Platform independence
(B) Separation of data from presentation
(C) Extensibility (ability to define custom tags)
(D) Built-in graphical rendering capabilities

223. Consider the following 2D geometric transformations:
(a) Translation
(b) Rotation
(c) Scaling
(d) Clipping
Which of the above are considered fundamental transformations?
(A) (a), (b), and (c) only
(B) (a) and (d) only
(C) (b), (c), and (d) only
(D) (a), (b), (c), and (d)

224. Consider the following statements about computer graphics concepts:
(a) A raster-scan system uses a frame buffer to store the image.
(b) A random-scan system draws images line by line from top to bottom.
(c) The mid-point circle algorithm avoids floating-point calculations.
(d) Polygon clipping is required to display only the parts of a polygon that lie within the viewport.
(e) The viewing pipeline transforms world coordinates into device coordinates.
Which of the above combination of statements is true?
(A) (a), (b), and (c) only
(B) (a), (c), (d), and (e) only
(C) (b), (d), and (e) only
(D) (a), (b), (d), and (e) only

225. What is a viewport in computer graphics?
(A) A region in world coordinates selected for display.
(B) A region on the display device where the window is mapped.
(C) The 3D volume that is visible to the camera.
(D) The coordinate system of a single object.

226. In the context of 3D viewing, what does a general projection transform do?
(A) It transforms a 3D view volume into a canonical view volume, typically a cube.
(B) It models the interaction of light with surfaces.
(C) It combines several objects into a single complex object.
(D) It defines the surface properties like color and texture.

227. What is the primary difference between Gouraud shading and Phong shading?
(A) Gouraud shading is a line drawing algorithm, while Phong shading is a polygon filling algorithm.
(B) Gouraud shading works only for 2D, while Phong shading is for 3D.
(C) Gouraud shading interpolates color intensities, while Phong shading interpolates surface normal vectors.
(D) Gouraud shading is computationally more expensive than Phong shading.

228. Consider the following statements regarding constructor calls in an inheritance hierarchy:
(a) Derived class constructor first, then base class constructor.
(b) Base class constructor first, then derived class constructor.
(c) Only the derived class constructor is called.
(d) Only the base class constructor is called.
(e) The order is determined at runtime.
Which of the above represents the correct order of constructor calls when a derived class object is created?
(A) (a)
(B) (b)
(C) (c)
(D) (e)

229. Which of the following is NOT a stage in the typical translation (compilation) process?
(A) Lexical Analysis
(B) Syntax Analysis
(C) Debugging

Click Here to Buy book on Amazon

213. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.

Assertion (A) : In C, strings are represented as null-terminated arrays of characters.

Reason (R) : The standard string library functions (like `strcpy`, `strlen`) rely on the null character (`\0`) to determine the end of the string.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

214. Match the following C++ keywords or concepts in List I with their purposes in List II.

List I: (C++ Keyword/Concept)
P. static (in a class)
Q. const (for a member function)
R. new
S. delete

List II: (Purpose)
1. Allocates memory for an object on the heap.
2. Creates a member that is shared by all objects of the class.
3. Deallocates memory previously allocated on the heap.
4. Indicates that a member function will not modify the object's data members.

(A) P-2, Q-4, R-1, S-3
(B) P-1, Q-2, R-3, S-4
(C) P-2, Q-1, R-4, S-3
(D) P-4, Q-2, R-1, S-3

215. Consider the following C preprocessor directives:
(a) `#include`
(b) `#define`
(c) `#ifdef`
(d) `#virtual`

Which of the above are valid preprocessor directives?
(A) (a) and (b) only
(B) (a), (b) and (c) only
(C) (c) and (d) only
(D) (a), (b), (c) and (d)

216. Identify the incorrect statement about Object-Oriented Programming.
(A) Polymorphism means one name, one form.
(B) Encapsulation combines data and functions into a single unit.
(C) Inheritance promotes code reusability.
(D) An abstract class cannot be used to create objects directly.

217. In C++, which stream is typically used for unbuffered error output?
(A) `cout`
(B) `cin`
(C) `cerr`
(D) `clog`

218. What is the homogeneous matrix representation for a 2D translation by (tx, ty)?
(A) `[[1, 0, tx], [0, 1, ty], [0, 0, 1]]`
(B) `[[1, 0, 0], [0, 1, 0], [tx, ty, 1]]`
(C) `[[tx, 0, 0], [0, ty, 0], [0, 0, 1]]`
(D) `[[1, tx, 0], [ty, 1, 0], [0, 0, 1]]`

219. Consider the following mechanisms in C++ related to file I/O and error handling:
(a) `try`/`catch` blocks for handling I/O errors.
(b) `ofstream` and `ifstream` for file stream operations.
(c) Virtual functions to read/write different file types.
(d) Templates for generic file formats.

Which of the above combination is most directly related to basic file handling and error management?
(A) (a) and (b) only
(B) (c) and (d) only
(C) (a) and (c) only
(D) (b) and (d) only

NODIA

220. A programmer needs to store a collection of items where each item has a different data type, but memory efficiency is critical, and only one item needs to be accessed at any given time. Which C construct is most suitable?
(A) `struct`
(B) `union`
(C) `array`
(D) `enum`

221. Which statement about C++ inheritance is incorrect?

cannot be changed through `p`.

Statement II : `int * const p;` declares a constant pointer `p` to an integer. The pointer `p` cannot be changed to point to another address.

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

239. What is a token in the C language?
(A) The smallest unit of a program, such as a keyword, operator, or identifier.
(B) A comment added by the programmer.
(C) An executable instruction.
(D) A header file included in the program.

240. Consider the following statements about C++ pointers and members:
(a) The `this` pointer holds the address of the object that invokes a member function.
(b) A static data member is shared among all objects of a class.
(c) A constructor can be declared `virtual`.
(d) A destructor should be declared `virtual` in a base class if it is to be inherited.

Which of the above statements are true?
(A) (a), (b), and (d) only
(B) (a) and (c) only
(C) (c) and (d) only
(D) (a), (b), (c), and (d)

241. What is the role of a line drawing algorithm like DDA or Bresenham's?
(A) To clip a line against a window.
(B) To determine which pixels on a raster display should be illuminated to represent a line segment.
(C) To transform a line from world coordinates to screen coordinates.
(D) To apply color to a line segment.

242. In a 2D shear transformation, if only the `shy` parameter is non-zero, what is the effect?
(A) It shifts points vertically, proportional to their x-coordinate.
(B) It shifts points horizontally.
(C) It scales the object uniformly.
(D) It rotates the object.

243. Consider the following representation techniques for complex 3D surfaces:
(a) Polygon surfaces using a low number of polygons.
(b) Quadric surfaces.
(c) Spline surfaces (like Bezier or B-spline surfaces).
(d) A single implicit equation for the entire car.
(e) Voxel representation.

Which of the following is considered the standard for industry design for smooth curved surfaces?
(A) (a)
(B) (b)
(C) (c)
(D) (d)

244. Identify the correct pair.
(A) `fopen` : A C++ stream operator
(B) `argv` : A count of command line arguments in C
(C) `struct` : A data type where all members share memory
(D) `polymorphism` : The ability to take more than one form

245. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.

Assertion (A) : Templates in C++ are a mechanism to implement generic programming.

Reason (R) : Templates are expanded by the compiler at runtime, which allows for dynamic type checking.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

246. What is the primary function of the Cohen-Sutherland line clipping algorithm?
(A) To draw a line using only integer arithmetic.
(B) To efficiently accept or reject a line segment against a rectangular clipping window.
(C) To fill a polygon with a solid color.
(D) To apply a perspective projection to a line segment.

NODIA

(D) (b) and (e)

180. Match the following language paradigms in List I with representative languages in List II.
List I: (Language Paradigm)
P. Procedural
Q. Object-Oriented
R. Functional
S. Logic
List II: (Representative Language)
1. LISP
2. Prolog
3. C
4. C++
(A) P-3, Q-4, R-1, S-2
(B) P-3, Q-4, R-2, S-1
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-1, R-3, S-4

181. Which web technology is a platform-independent, server-side component for extending the functionality of a web server?
(A) HTML
(B) DHTML
(C) Java Servlet
(D) Java Applet

182. What is the key difference between the boundary-fill and flood-fill algorithms?
(A) Boundary-fill checks for a boundary color, while flood-fill checks for an interior color to replace.
(B) Flood-fill checks for an interior color to replace, while boundary-fill checks for a boundary color.
(C) Boundary-fill works on any shape, while flood-fill only works on rectangles.
(D) Flood-fill is faster but uses more memory.

183. Consider the following statements about geometric transformations:
(a) A rigid body transformation preserves distances and angles (e.g., translation, rotation).
(b) Concatenating transformation matrices is generally more efficient than applying each transformation separately.
(c) Transformation matrices are not commutative, i.e., T1 * T2 is not always equal to T2 * T1.
(d) A uniform scaling transformation preserves the shape of an object.
Which of the above statements are true?
(A) (a) and (b) only
(B) (a), (c) and (d) only
(C) (c) and (d) only
(D) (a), (b), (c) and (d)

184. The following question consists of two statements, one labeled as Assertion (A) and the other as Reason (R). Examine these two statements carefully and select the correct answer using the codes given below.
Assertion (A) : The `const` keyword in C++ can be used to declare variables whose values cannot be changed after initialization.
Reason (R) : Using `const` helps in writing safer code by preventing accidental modification of variables that are intended to be constant, and it allows the compiler to perform certain optimizations.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

185. What is the role of the `Viewing Funnel` or `View Frustum` in 3D viewing?
(A) It is the coordinate system of the display device.
(B) It is a method for representing curved surfaces.
(C) It defines the region of 3D space that is visible in the final image.
(D) It is an algorithm for smooth shading of polygons.

186. Which of the following is NOT a type of spline representation?
(A) Polygon Spline
(B) Bezier Spline
(C) B-Spline
(D) Cardinal Spline

187. What does orthogonality in programming language design refer to?
(A) Capability of source code to execute across multiple OS.
(B) Strictness of the internal data types within the syntax.

Downloaded from www.pdf.tube

251. In the context of language translation, what is the primary role of a virtual computer?
(A) To execute the final machine code directly on the hardware.
(B) To offer an abstract, consistent target machine, simplifying the translation process.
(C) To provide a suite of development tools for programmers.
(D) To serve as a concrete hardware specification for a new processor.

252. Which of the following correctly lists distinct stages in the language translation process?
(A) Lexical analysis, syntax analysis, semantic analysis, and code generation.
(B) Binding, execution, debugging, and compilation.
(C) Object-oriented design, procedural implementation, and functional testing.
(D) Syntax definition, environment setup, and virtual machine configuration.

253. The association of a variable with its data type at compile time is an example of which concept?
(A) Programming paradigm
(B) Formal transition model
(C) Binding time
(D) Programming environment

254. Based on the paragraph, how are programming language concepts and paradigms related?
(A) Paradigms are tools within a programming environment, while concepts are theoretical.
(B) Paradigms and concepts are interchangeable terms for the syntax of a language.
(C) Concepts are derived from formal transition models, whereas paradigms define the virtual computer.
(D) Concepts are the fundamental ideas that are organized into broader programming paradigms.

255. Formal transition models, such as finite automata, are most directly applied to which stage of translation?
(A) Semantic analysis
(B) Code optimization
(C) Lexical analysis
(D) Intermediate code generation

Direction : Please read the following source carefully.

In programming languages, data types are fundamental constructs that define the nature of the values a variable can hold and the operations that can be performed on those values. Every data type is characterized by a set of properties, which dictate its storage requirements, representation, and the domain of permissible values. The concept of an object, in this context, refers to a runtime instance of a data type, which occupies a region of memory and has a specific value. A critical distinction is made between elementary data types, which are generally atomic and indivisible from the programmer's perspective. These are broadly categorized into scalar and composite types. Scalar data types represent a single value, such as an integer, a floating-point number, a character, or a Boolean value. They are the basic building blocks for all other data structures. In contrast, composite data types are constructed from one or more elementary types. They represent collections of data and include structures like arrays, which are homogeneous collections of elements, and records (or structures in C), which are heterogeneous collections of fields. The properties of a composite type are derived from the properties of its constituent scalar or other composite types, along with the rules governing its structure and access methods. Understanding this hierarchy is essential for effective data modeling and memory management in any programming language.

Based on the above source, attempt the following five questions.

NODIA

256. According to the passage, what is the primary characteristic of a scalar data type?
(A) It is constructed from multiple other data types.
(B) It represents a single, indivisible value.
(C) It represents a collection of heterogeneous data.
(D) It is a runtime instance occupying memory.

257. Which of the following best defines a composite data type?
(A) A type that can only hold a single integer or character value.
(B) A theoretical concept with no runtime instance or memory allocation.

ANSWERS

1. (B) A fundamental style or model of computer programming.
Computer programming paradigms serve as the foundational frameworks or styles that dictate how a programmer organizes and structures code. These paradigms provide distinct approaches to problem-solving, such as procedural, object-oriented, or functional models. By following a specific model, developers adhere to standardized principles for designing logic and managing data effectively within software systems.

2. (A) The moment a name is linked to an object or property.
Binding time represents the specific instance during the program lifecycle when an association is established between an identifier and its corresponding property. This process occurs at various stages, including design, compile, link, or execution phases. Static binding happens before the program runs, whereas dynamic binding occurs during the actual execution, allowing for greater flexibility and runtime behavior.

3. (C) A character
Scalar data types are the most basic units of data representation that hold a single value at any given time. These types are atomic, meaning they cannot be decomposed into smaller constituent parts. Common examples include integers, floating-point numbers, and characters. Unlike composite types like arrays or structures, scalar types represent individual data points without internal organizational complexity.

4. (A) Both A and R are true and R is the correct explanation of A.
Static binding enhances program efficiency because the association between identifiers and memory addresses or subprograms is resolved during compilation. Since the compiler knows the exact location of the resources beforehand, it generates machine code that executes direct calls or accesses. This eliminates the need for expensive runtime lookups or indirect addressing, resulting in faster and optimized execution.

5. (C) Lexical Analysis
Lexical analysis represents the initial phase of the compiler's translation process, focusing on reading the source code as a stream of characters. It systematically groups these individual characters into meaningful sequences known as lexemes. These lexemes are then categorized into symbolic representations called tokens, such as keywords, identifiers, or operators, which serve as input for subsequent syntax analysis.

6. (B) An object is an instance of a class.
In object-oriented programming, a class functions as a blueprint or template that defines the structure and behavior for a set of entities. An object is a concrete, tangible instance created based on that class definition. While the class specifies the attributes and methods available, the object holds specific data values and occupies memory during the program's actual execution.

7. (D) To achieve runtime polymorphism.
Virtual functions in C++ facilitate runtime polymorphism by allowing a base class to define methods that derived classes can override with specific implementations. When a virtual function is invoked through a base class pointer or reference, the program determines the actual type of the object at execution time and executes the appropriate version, ensuring flexible and dynamic behavior.

8. (A) Random-scan display
Random-scan displays, often referred to as vector-refresh systems, operate by directing the electron beam exclusively to the coordinates where the image components are located. Instead of scanning the entire screen line by line, the beam traces the geometric primitives directly. This targeted approach results in smooth lines and is particularly effective for high-resolution line-drawing applications in computer graphics.

9. (A) P-2, Q-4, R-1, S-3
C language concepts include pointers, which store memory addresses, and structures that group variables of different types. Unions

Click Here to Buy book on Amazon

the same memory location.
(D) It represents the smallest lexical unit, such as a keyword or an operator.

264. The mechanism in C that facilitates passing parameters to a program at the time of its execution is known as:
(A) File handling
(B) Subprogram control
(C) Preprocessor directives
(D) Command line arguments

265. Which of the following are considered tokens in the C language?
(A) Identifiers, keywords, and operators.
(B) Arrays, structures, and unions.
(C) Pointers, functions, and files.
(D) Subprogram control and sequence control.

Direction : Please read the following source carefully.

Object-Oriented Programming (OOP) is a paradigm centered around the concept of objects, which encapsulate both data and the procedures that operate on that data. The blueprint for creating objects is a class, which is a user-defined type that bundles data members (attributes) and member functions (methods) into a single unit. The process of creating an object from a class is known as instantiation. A core principle of OOP is encapsulation, which is the mechanism of restricting direct access to some of an object's components, effectively hiding the internal state and requiring interaction through a defined public interface. This principle is fundamental to data protection and modularity. Another key concept is inheritance, which allows a new class (a derived or subclass) to acquire the properties and behaviors of an existing class (a base or superclass). This promotes code reuse and establishes a hierarchical relationship between classes. Polymorphism, meaning many forms, enables objects of different classes to be treated as objects of a common superclass, allowing a single function or operator to behave differently depending on the object it is acting upon. Finally, an abstract class is a special type of class that cannot be instantiated on its own and is designed to be inherited by other classes. It often defines an interface with one or more abstract methods that derived classes must implement, enforcing a certain structure and functionality.

Based on the above source, attempt the following five questions.

266. Which OOP concept refers to the creation of an object from a class?
(A) Encapsulation
(B) Instantiation
(C) Inheritance
(D) Polymorphism

267. The principle of bundling data and methods within a single unit and restricting direct access to the data is known as:
(A) Polymorphism
(B) Encapsulation
(C) Inheritance
(D) Instantiation

268. What is the primary purpose of an abstract class in Object-Oriented Programming?
(A) To serve as a base class that cannot be instantiated but defines an interface for subclasses.
(B) To be directly instantiated multiple times for creating simple objects.
(C) To encapsulate data without any associated methods or behaviors.
(D) To allow a single function to behave differently based on the object's runtime type.

269. Inheritance is a mechanism that primarily promotes:
(A) Data hiding and protection of an object's internal state.
(B) The ability of a single function to operate on objects of different classes.
(C) Code reuse by allowing a new class to acquire properties from an existing class.
(D) The creation of multiple object instances from a single class blueprint.

NODIA

270. A programmer wants to define a common interface for a group of related classes, ensuring that each class implements a specific set of methods. Which OOP feature is most suitable for this purpose?
(A) Encapsulation
(B) Instantiation
(C) A standard class object

specified boundary color or a new color region is encountered.

Based on the above source, attempt the following five questions.

281. What is the fundamental difference between a Raster-Scan and a Random-Scan system?
(A) Raster-scan sweeps the entire screen row-by-row, while random-scan directs the beam only where the image is.
(B) Raster-scan uses an electron beam, while random-scan uses input devices.
(C) Raster-scan draws pictures one line at a time, while random-scan draws the entire screen at once.
(D) Raster-scan is for 3D objects, while random-scan is for 2D primitives like points and lines.

282. The primary purpose of line drawing algorithms like Bresenham's is to:
(A) Fill a polygon with color by starting from an interior point.
(B) Determine which pixels to illuminate to approximate a line on a raster display.
(C) Draw a line by directing an electron beam between two endpoints on a vector display.
(D) Transform, scale, and rotate a line in 2D space.

283. A graphics programmer needs to fill an irregularly shaped region defined by a single border color. Which algorithm would be most appropriate?
(A) Mid-Point Circle algorithm
(B) Boundary-Fill algorithm
(C) Scan Line Polygon Fill algorithm
(D) A line drawing algorithm

284. How does the Scan Line Polygon Fill algorithm operate?
(A) By processing the polygon's area on a horizontal, row-by-row basis.
(B) By starting from a seed point and recursively filling adjacent pixels.
(C) By tracing the boundary of the polygon and filling inwards.
(D) By directing an electron beam to trace the outline of the polygon.

285. According to the passage, which of the following is characteristic of a modern graphics monitor?
(A) It typically uses a Random-Scan system for high-precision drawing.
(B) It relies on Boundary-Fill algorithms for its primary display mechanism.
(C) It represents images as a grid of pixels and uses a Raster-Scan approach.
(D) It draws images by moving a pencil-like stylus across the screen.

Direction : Please read the following source carefully.

In 2-D computer graphics, geometrical transforms are fundamental operations used to manipulate the position, size, and orientation of objects. The primary transformations include translation (moving an object), scaling (resizing an object), and rotation (reorienting an object about a point). Other important transforms are reflection (creating a mirror image) and shear (distorting the shape of an object). These transformations can be mathematically represented using matrices. By representing 2-D points `(x, y)` as vectors, a 2x2 matrix can perform scaling, rotation, reflection, and shear centered at the origin. However, translation is an additive operation and cannot be expressed with a 2x2 matrix multiplication. To unify all transformations into a single matrix multiplication framework, homogeneous coordinates are introduced. A 2-D point `(x, y)` is represented as `(x, y, 1)` in homogeneous coordinates, and all transformations are represented by 3x3 matrices. This allows for the creation of composite transforms by simply multiplying the individual transformation matrices in sequence. The process of displaying a 2-D scene involves a viewing pipeline. This pipeline includes transformations between coordinate systems, notably the window-to-viewport coordinate transformation, which maps a selected part of the scene (the window) onto a specified area of the display device (the viewport). Viewing functions set up this mapping, and clipping algorithms (e.g., for lines and polygons) are used to discard parts of the objects that lie outside the viewing window.

Based on the above source, attempt the following five questions.

Downloaded from www.pdf.tube

XML is designed to describe and structure data. The client-side interactivity is primarily driven by scripting languages like JavaScript, which run in the user's browser to manipulate the page content in real-time. On the server side, more complex logic, database interaction, and content generation are handled by technologies like Java. Within the Java ecosystem for the web, Servlets are server-side components that receive requests and generate dynamic responses, running within a Java-enabled web server. Applets, on the other hand, are small Java applications designed to be downloaded from a web server and executed within a web browser, providing a way to embed complex client-side functionality directly into a web page.

Based on the above source, attempt the following five questions.

276. What is the primary functional difference between HTML and XML as described in the passage?
(A) HTML is focused on the presentation of content, while XML is focused on describing data.
(B) HTML is used for server-side logic, while XML is for client-side scripting.
(C) HTML is a scripting language, while XML is a compiled language.
(D) HTML is used for dynamic content, while XML is for static page structure.

277. According to the text, which technology is a server-side component designed to handle requests and generate dynamic responses?
(A) DHTML
(B) Applet
(C) XML
(D) Servlet

278. The passage describes DHTML as:
(A) A standalone programming language for server-side development.
(B) A combination of technologies used to dynamically modify web page content.
(C) A markup language for structuring data exchange.
(D) A small application that runs inside a web browser to provide complex functionality.

279. How does an Applet differ from a Servlet?
(A) Applets run on the server to process requests, while Servlets run in the client's browser.
(B) Applets are written in HTML, while Servlets are written in Java.
(C) Applets are executed within a web browser on the client-side, while Servlets run on the web server.
(D) Applets are for data structuring, while Servlets are for page presentation.

280. Client-side interactivity, such as validating a form before submission, is most directly achieved using which of the following?
(A) Servlets
(B) XML
(C) The core HTML specification
(D) Scripting

Direction : Please read the following source carefully.

Computer graphics involves the creation and manipulation of images using computers, relying on specialized hardware and algorithms. Video-display devices are the primary output medium, and they can be broadly classified based on their drawing methodology into Raster-Scan and Random-Scan systems. Raster-scan systems, common in graphics monitors today, represent an image as a grid of pixels and draw it by sweeping an electron beam across the entire screen, one row at a time. In contrast, Random-scan (or vector) systems draw an image by directing the electron beam only to the parts of the screen where the image is to be drawn, much like drawing with a pencil on paper. The fundamental graphical primitives are points and lines, which form the basis of more complex shapes. Efficiently drawing these primitives is crucial, leading to the development of specialized algorithms. Line drawing algorithms, such as Bresenham's, are designed to determine which pixels on a raster display should be illuminated to approximate a straight line. Similar logic extends to circles and ellipses with algorithms like the Mid-Point Circle algorithm. For filling complex shapes, polygon fill algorithms are used. The Scan Line Polygon Fill algorithm works by processing the polygon on a row-by-row basis, while Boundary-Fill and Flood-Fill algorithms work by starting at an interior point and spreading outwards until a

NODIA

large number of triangular facets.

293. According to the passage, a significant advantage of B-spline curves over Bezier curves is:
(A) They are defined by a simpler, second-degree equation.
(B) They always pass through every one of their control points.
(C) They are limited to representing only 2-D objects.
(D) They offer more localized control over the curve's shape.

294. A sphere or a cylinder would be most efficiently and accurately represented as a:
(A) Polygon surface
(B) Quadric surface
(C) Bezier surface
(D) B-spline curve

295. The extension of Bezier and B-spline curves to model smooth, complex object skins in three dimensions results in:
(A) Quadric surfaces
(B) Polygon rendering methods
(C) Bezier and B-spline surfaces
(D) General projection transforms

Direction : Please read the following source carefully.

The process of rendering a realistic 3-D scene involves a sequence of steps known as the viewing pipeline, which transforms 3-D object representations into a 2-D image on a display. This pipeline manages various coordinate systems and transformations. An essential part of creating realism is the application of illumination models, which simulate how light interacts with surfaces. These models calculate the color and brightness of points on a surface based on light sources, material properties, and viewing position. The results of these calculations are then used by polygon rendering methods to determine the final colors for the pixels covering the polygons. For example, methods like Gouraud and Phong shading compute colors across a polygon's surface to create a smooth appearance. The geometric aspect of the 3-D viewing pipeline involves general projection transforms, which project the 3-D world onto a 2-D view plane. These projections can be parallel or perspective, with perspective projections providing depth realism by making distant objects appear smaller. After projection, clipping is performed in 3-D space to remove objects or parts of objects that fall outside the defined view volume. This entire process, from object representation in 3-D coordinates to the final 2-D image, ensures that a three-dimensional model is correctly and realistically displayed from a specific viewpoint.

Based on the above source, attempt the following five questions.

296. What is the primary function of an illumination model in 3-D graphics?
(A) To project a 3-D object onto a 2-D view plane.
(B) To remove parts of an object that are outside the view volume.
(C) To represent a 3-D object using polygon surfaces or splines.
(D) To simulate the interaction of light with surfaces to determine their color and brightness.

297. The process of making distant objects appear smaller to create depth realism is a characteristic of which type of transformation?
(A) Perspective projection
(B) Parallel projection
(C) Polygon rendering
(D) Illumination modeling

298. Polygon rendering methods like Gouraud and Phong shading are used to:
(A) Define the 3-D coordinates of an object.
(B) Compute colors across a polygon's surface for a smooth appearance.
(C) Clip polygons against the view volume.
(D) Transform objects from world coordinates to view coordinates.

299. Within the 3-D viewing pipeline, what is the purpose of clipping?
(A) To apply an illumination model to a surface.
(B) To remove parts of objects that lie outside the defined view volume.
(C) To make distant objects appear smaller.
(D) To transform a 3-D representation into a 2-D image.

300. The entire sequence of operations that transforms 3-D object representations into a 2-D

Click Here to Buy book on Amazon

286. What is the primary advantage of using homogeneous coordinates in 2-D graphics?
(A) They reduce the memory required to store points and vectors.
(B) They are only used for reflection and shear transformations.
(C) They simplify the process of clipping polygons against the viewport.
(D) They allow all geometric transformations, including translation, to be represented as matrix multiplications.

287. A composite transform that first rotates an object and then scales it is created by:
(A) Adding the rotation and scaling matrices.
(B) Averaging the elements of the rotation and scaling matrices.
(C) Multiplying the individual transformation matrices for rotation and scaling.
(D) Applying a viewing function that combines rotation and scaling.

288. The window-to-viewport coordinate transformation is responsible for:
(A) Distorting the shape of an object using a shear operation.
(B) Mapping a selected part of the scene onto an area of the display device.
(C) Removing the parts of an object that lie outside a specified boundary.
(D) Representing a 2-D point with three coordinates for matrix multiplication.

289. Which of the following transformations is described as distorting the shape of an object, for example, by slanting it?
(A) Translation
(B) Rotation
(C) Reflection
(D) Shear

290. In the context of the 2-D viewing pipeline, what is the role of a clipping algorithm?
(A) To discard portions of objects that are outside the viewing window.
(B) To create a mirror image of an object.
(C) To combine multiple transformation matrices into one.
(D) To convert a 2-D point into homogeneous coordinates.

Direction : Please read the following source carefully.

Representing three-dimensional objects is a core task in computer graphics, with various schemes balancing realism, complexity, and computational cost. A common and straightforward method is to use polygon surfaces, where the object's boundary is approximated by a mesh of polygonal facets, such as triangles or quadrilaterals. For objects with regular curved surfaces, quadric surfaces, which are defined by second-degree equations (e.g., spheres, cylinders, and cones), offer a more concise and accurate representation than a polygonal mesh. For modeling more complex, free-form shapes with smooth curves and surfaces, spline representations are employed. Splines are piecewise polynomial functions that provide a high degree of continuity and local control. Prominent examples in this category are Bezier and B-spline curves. A Bezier curve is defined by a set of control points that influence its shape, though the curve does not necessarily pass through all of them. B-spline curves offer more localized control, meaning that moving one control point only affects a small portion of the curve. These curve formulations can be used into three dimensions to create Bezier and B-spline surfaces, which are defined by a mesh of control points and are widely used in computer-aided design for creating smooth, complex object skins like car bodies or airplane fuselages.

Based on the above source, attempt the following five questions.

291. Which representation scheme is most suitable for approximating the boundary of a complex 3-D object with a mesh of flat facets?
(A) Quadric surfaces
(B) B-spline curves
(C) Polygon surfaces
(D) Bezier curves

292. For what purpose are Bezier and B-spline representations primarily used?
(A) To model complex, free-form shapes with smooth curves and surfaces.
(B) To represent objects with regular, mathematically simple surfaces like spheres.
(C) To define the viewing pipeline for 3-D projection.
(D) To approximate an object's surface with a

Downloaded from www.pdf.tube

(D) An abstract class

Direction : Please read the following source carefully.

C++ extends the C language with object-oriented features and other enhancements. Like C, its syntax is built upon tokens such as identifiers, variables, constants, and operators. It supports a comprehensive set of data types and control statements for managing program flow. A significant aspect of C++ is its advanced function handling, including parameter passing mechanisms and the concept of virtual functions, which are crucial for implementing runtime polymorphism. The core of its object-oriented nature lies in the class and object constructs. Classes define blueprints for objects, and their lifecycle is managed by special member functions: constructors, which initialize objects upon creation, and destructors, which clean up resources when objects are destroyed. C++ supports extensive overloading, allowing functions and operators to have multiple implementations depending on their context. Inheritance allows for building class hierarchies, while templates provide a powerful mechanism for generic programming, enabling the creation of functions and classes that work with any data type. For robustness, C++ includes sophisticated exception and event handling mechanisms to manage runtime errors gracefully. The language also provides extensive support for streams and files for I/O operations and facilitates the development of multifile programs, where different parts of the code are organized into separate files for better modularity and management.

Based on the above source, attempt the following five questions.

271. In C++, what is the primary role of a virtual function?
(A) To initialize an object when it is created.
(B) To enable the creation of classes that work with any data type.
(C) To implement runtime polymorphism within a class hierarchy.
(D) To handle runtime errors and exceptional events.

272. Which C++ feature is specifically designed to manage the initialization and cleanup of objects?
(A) Overloading and templates
(B) Constructors and destructors
(C) Virtual functions and inheritance
(D) Exception handling and streams

273. The ability to write a single function or class that can operate on different data types is achieved in C++ through the use of:
(A) Inheritance
(B) Virtual functions
(C) Overloading
(D) Templates

274. Based on the passage, how does C++ facilitate the development of large, modular applications?
(A) By supporting the organization of code into multifile programs.
(B) By limiting programs to a single source file.
(C) By removing support for control statements and operators.
(D) By relying solely on automatic garbage collection.

275. The mechanism in C++ that allows a function or an operator to have multiple implementations based on the types or number of its arguments is called:
(A) Inheritance
(B) Polymorphism
(C) Overloading
(D) Encapsulation

NODIA

Direction : Please read the following source carefully.

Web programming encompasses a variety of technologies designed to create dynamic and interactive content for the World Wide Web. At its core is HTML (HyperText Markup Language), which provides the fundamental structure for web pages. To enhance the presentation and interactivity of these pages, DHTML (Dynamic HTML) is used, which is not a separate language but a combination of HTML, CSS, and scripting to enable dynamic modification of a page's content after it has been loaded. For data representation and exchange between disparate systems, XML (eXtensible Markup Language) is employed; unlike HTML, which is focused on presentation,

image, including projections and clipping, is known as the:

(A) Viewing pipeline
(B) Illumination model
(C) Polygon rendering method
(D) General projection transform

www.nodia.in

Click Here to Buy book on Amazon

(C) A type defined by the domain of permissible operations but not values.
(D) A type constructed from one or more elementary data types to form a collection.

258. In the context of the passage, what is the relationship between a data type and an object?
(A) An object is a runtime instance of a data type.
(B) A data type is a runtime instance of an object.
(C) Data types and objects are identical concepts.
(D) An object defines storage, while a data type defines operations.

259. Based on the description, an array of integers would be classified as what kind of data type?
(A) Scalar, because it represents numbers.
(B) Object, because it has properties.
(C) Composite, because it is a collection of elements.
(D) Atomic, because its elements are indivisible.

260. The passage suggests that the properties of a composite type, such as a structure in C, are determined by:
(A) The properties of its constituent types and its structural rules.
(B) The operating system's memory manager exclusively.
(C) The number of operations that can be performed on it.
(D) A single scalar value it is designed to represent.

Direction : Please read the following source carefully.

The C programming language provides a rich set of features for procedural programming, emphasizing low-level memory access and system control. The fundamental lexical units of a C program are tokens, which include identifiers, keywords, constants, operators, and other separators. Identifiers are names given to entities like variables, functions, and user-defined data types. C supports a variety of data types, both scalar (like `int`, `char`) and composite (like arrays, structures, and unions). Program flow is managed through sequence control mechanisms, such as conditional statements and loops, and subprogram control, which involves the definition and invocation of functions. A key feature of C is its extensive use of pointers, which are variables that store memory addresses, enabling direct memory manipulation and efficient handling of complex data structures. Arrays, as homogeneous collections, and structures, as heterogeneous collections, allow for the organization of related data. The union is a special data type that allows different data types to be stored in the same memory location. Strings are managed as null-terminated arrays of characters. C also provides robust file handling capabilities for input and output operations with secondary storage. Additionally, the preprocessor allows for macro expansion and file inclusion before actual compilation, while command line arguments provide a mechanism for passing parameters to a program upon execution.

Based on the above source, attempt the following five questions.

261. What is the primary function of a pointer in the C programming language?
(A) To define a user-defined data type with heterogeneous fields.
(B) To execute pre-compilation tasks like macro expansion.
(C) To manage the sequential flow of control within a program.
(D) To store a memory address, enabling direct memory manipulation.

262. According to the passage, how does a structure differ from an array in C?
(A) Structures are for subprogram control, while arrays are for sequence control.
(B) Structures store heterogeneous data, while arrays store homogeneous data.
(C) Structures are handled by the preprocessor, while arrays are handled by the compiler.
(D) Structures can only store scalar types, while arrays can store composite types.

263. What is the defining characteristic of a union data type as described in the text?
(A) It combines multiple functions into a single subprogram.
(B) It is a collection of elements of the same data type.
(C) It allows different data types to be stored in

are unique because they allow different data members to share the same memory space, minimizing total storage. Preprocessor directives provide instructions that are handled before the compilation phase begins. These elements combined offer powerful tools for memory management and code organization.

10. (A) array
Composite data types are constructed by combining multiple elementary or other composite types into a single structure. An array is a primary example because it organizes a collection of elements of the same type under one name, accessible via indices. This contrasts with scalar types like integers or booleans, which store only one individual piece of information at once.

11. (D) hardware abstraction providing a language interface.
A virtual computer in language design refers to the software layer that provides a uniform interface to the programmer, hiding the complexities of the underlying physical hardware. It creates an abstract execution environment where a programming language's instructions can be processed consistently regardless of the specific machine architecture, simplifying the development process and enhancing portability across different systems.

12. (C) Statement I is false, but Statement II is true.
Pointers in C are specialized variables that store the memory addresses of other variables rather than the values themselves. To access the actual data stored at that address, the dereference operator is used. Additionally, the address-of operator retrieves the specific location of a variable in memory. These operators are fundamental for direct memory manipulation and efficient data handling.

13. (C) It distorts the shape of an object by sliding layers over each other.
Shearing is a two-dimensional geometric transformation that distorts the shape of an object by shifting its coordinates along a specific axis. This effect is analogous to sliding layers of an object over one another, causing a rectangular shape to appear slanted or skewed. Unlike scaling or rotation, shearing modifies the object's proportions by changing its angular relationships without resizing.

14. (D) (a), (b), (c) and (d)
Programming paradigms offer diverse models for structuring computation. Procedural programming focuses on procedure calls, while object-oriented models organize logic around objects. Functional programming emphasizes the evaluation of mathematical functions, avoiding state changes. Logic programming uses formal logic to define relationships and rules. Each paradigm provides a different perspective on problem-solving, and many modern languages integrate features from several models.

15. (A) Extensible Markup Language
XML represents Extensible Markup Language, a versatile tool designed for storing and transporting data in a structured format. Unlike HTML, which focuses on the presentation of information, XML emphasizes the actual content by allowing users to define their own custom tags. This flexibility makes it a standard choice for data exchange between different systems, platforms, and various software applications.

16. (B) The memory address of the object.
A data type is fundamentally defined by the set of values it can hold, the operations that can be performed on those values, and its internal memory representation. While every object in a running program occupies a specific memory address, that address is an attribute of the instance rather than a fundamental property of the data type itself.

17. (B) To allow different variables to share the same memory location.
The union keyword in C is used to define a data structure where all member variables share the same memory space. Only one member can store a value at any given time, and the size of the union is determined by its largest member. This mechanism is primarily utilized to conserve memory when multiple mutually exclusive data elements are needed.

18. (A) Cannot be used to create objects.
An abstract class in C++ serves as a foundational blueprint for other classes and cannot be instantiated directly. It typically contains at least

Downloaded from www.pdf.tube

74. (A) Both A and R are true and R is the correct explanation of A.
Abstract classes are designed to act as general templates for a class hierarchy and cannot be used to create objects directly. They must contain at least one pure virtual function, which acts as a placeholder for functionality that subclasses must implement. This ensures that the abstract class remains incomplete while providing a consistent interface and structural foundation for all derived classes.

75. (D) my_func(&x);
To allow a function to modify a variable's value, the address of that variable must be passed as an argument. This is known as call-by-reference using pointers. By passing the address of an integer variable to a function that accepts an integer pointer, the function can directly access and alter the original variable's memory, ensuring that changes persist after the function returns.

76. (B) To create functions and classes that can operate with generic types.
The template keyword in C++ enables generic programming by allowing functions and classes to be defined with placeholder data types. This allows a single implementation to work with various types, such as integers, floats, or custom objects. The compiler generates the specific code for each type used, which promotes code reuse and reduces duplication while maintaining strong type safety during compilation.

77. (C) Finite Automaton
A finite automaton is a formal mathematical model used to describe systems with a finite number of states and transitions between them. In language theory, these models are essential for defining and recognizing regular languages. They are widely used in the design of lexical analyzers to identify tokens within source code, providing a rigorous framework for understanding computational processes and language structures.

78. (D) Checking for type mismatches and other meaning-related errors.
Semantic analysis is the stage of compilation where the meaning of the program is verified. It ensures that the code follows the language's logical rules, such as verifying that variables are declared before use and that operations are performed on compatible data types. This phase detects errors that syntax analysis cannot, ensuring the program's overall consistency and logical correctness before code generation.

79. (B) P-3, Q-4, R-1, S-2
Web technologies have specialized roles in creating and managing online content. HTML defines the structure and content of pages, while XML is used for platform-independent data storage and transport. Java Servlets handle server-side processing to generate dynamic responses. DHTML combines multiple technologies to create interactive and animated user experiences, allowing web pages to change dynamically in the browser without reloading.

80. (B) Attempting to use a reference variable that has not been initialized to an object.
A NullPointerException in Java typically occurs when the program attempts to access or modify an object through a reference variable that currently points to nothing. This happens when a variable is declared but not assigned to a valid object instance. Since there is no actual object in memory to perform the requested operation, the Java Runtime Environment throws this exception.

81. (B) Convex hull property
The convex hull property ensures that a Bezier curve is always contained within the smallest convex shape that encloses all its control points. This provides a useful geometric bound for the curve, which is valuable for collision detection and clipping in computer graphics. It guarantees that the curve will not stray far from the path defined by its control points.

82. (A) Structures allocate memory for all members; unions allocate only for the largest.
The fundamental difference between C structures and unions lies in their memory management. A structure allocates separate memory for each of its members, allowing all of them to be accessed simultaneously. In contrast, a union allocates a single shared memory block large enough for its biggest member, meaning only one member can be reliably accessed at any given time to save space.

NODIA

28. (B) Scan Line Polygon Fill Algorithm
The scan-line polygon fill algorithm is designed for efficiency by processing a polygon one horizontal line at a time. It identifies the intersections between the scan line and the polygon's edges, then fills the spans between interior intersection points. This systematic approach is more suitable for complex shapes and large areas compared to recursive methods like boundary-fill or flood-fill.

29. (A) Both A and R are true and R is the correct explanation of A.
Declaring a function as virtual in a base class is the fundamental requirement for achieving runtime polymorphism in C++. This mechanism ensures that the function call is resolved at execution time based on the actual object type. Consequently, when using a base pointer, the program invokes the specific implementation provided by the derived class, allowing for dynamic behavior.

30. (D) <a>
The <a> tag, or anchor tag, is the standard HTML element used to create hyperlinks between web pages or to specific locations within a page. It uses the href attribute to specify the destination URL. When a user clicks on the content enclosed within the anchor tags, the browser navigates to the linked address, facilitating web navigation and connectivity.

31. (A) A surface described by a second-degree equation (e.g., spheres, ellipsoids).
Quadric surfaces are three-dimensional shapes defined by a general second-degree algebraic equation in three variables. Common examples include spheres, ellipsoids, cylinders, and cones. In computer graphics, these surfaces are valued for their mathematical precision and are often used to represent smooth, curved objects in geometric modeling and rendering, providing an alternative to approximate polygon-based meshes.

32. (B) Determining the order in which statements are executed.
Sequence control refers to the mechanisms in a programming language that govern the flow of execution within a program. This includes the linear progression of statements as well as control structures like loops, branches, and function calls. By managing the order of operations, sequence control ensures that computations are performed logically and that the program achieves its intended functional goals.

33. (B) It is a standalone language separate from HTML.
Dynamic HTML (DHTML) is not a single, distinct language but rather a combination of existing technologies used to create interactive websites. It integrates standard HTML for content, Cascading Style Sheets (CSS) for styling, and scripting languages like JavaScript for behavior. Together, these tools allow developers to modify web page content and appearance dynamically after the page has finished loading.

34. (B) P-3, Q-4, R-1, S-2
C++ utilizes specific keywords and syntax to implement its core object-oriented features. Inheritance is established using a colon in the class declaration, while object instantiation often involves the new operator for dynamic memory. Exception handling is managed through try and catch blocks. Runtime polymorphism is achieved by declaring member functions as virtual, allowing for dynamic method resolution during execution.

35. (C) Viewing Transformation
Viewing transformation is the process in computer graphics that maps objects from world coordinates to the specific display area of an output device. This involves defining a viewing window in the world and a corresponding viewport on the screen. The transformation ensures that the selected portion of the scene is correctly scaled and positioned within the viewport for the user.

36. (C) Statement I is true, but Statement II is false.
Operator overloading in C++ allows programmers to define custom behaviors for standard operators when applied to user-defined types like classes. However, this capability has limitations; certain operators, such as the scope resolution operator, member access operator, and the conditional operator, cannot be overloaded. These restrictions maintain the fundamental syntax and logical consistency of the C++ language.

37. (C) (x, y, 1)

Click Here to Buy book on Amazon

92. (B) A polygon surface must always be convex.
Polygon surfaces are a fundamental way to represent 3D objects by using a mesh of flat facets. While convex polygons are often preferred because they are easier to process and render, polygon surfaces can also consist of concave polygons. These surfaces are defined by vertex coordinates and attributes like color or texture, and they are rendered through various filling and shading techniques.

93. (C) It provides local control over the surface shape.
A B-spline surface is a powerful tool for modeling smooth 3D shapes because it offers local control. This means that modifying a single control point only affects a specific, localized area of the surface, rather than its entire shape. This characteristic makes B-spline surfaces highly desirable in design and engineering, as it allows for precise adjustments without unintended global consequences.

94. (A) Both A and R are true and R is the correct explanation of A.
The virtual keyword is the essential mechanism for enabling runtime polymorphism in C++. By declaring a function as virtual in the base class, the programmer tells the compiler to use dynamic binding. This ensures that when a base pointer refers to a derived object, the system correctly identifies and executes the derived class's version of the function, allowing for flexible behavior.

95. (A) The smallest individual unit of a program, such as a keyword, identifier, or operator.
In the context of programming languages, a token is the most basic lexical unit that carries meaning for the compiler or interpreter. During the lexical analysis phase, the source code is broken down into tokens, which include keywords, variables, operators, and literals. These tokens form the building blocks for the syntax and structure of the entire program, much like words in a sentence.

96. (B) #ifdef
The #ifdef directive is a C preprocessor command used for conditional compilation. it checks if a specific macro has been defined before the compilation of a code block. This allows programmers to include or exclude portions of code based on specific conditions, such as different operating systems or debugging modes, providing greater flexibility and control over the final executable's features and performance.

97. (B) Call-by-value passes a copy of the argument's value, while call-by-reference passes the argument itself, allowing modification.
Call-by-value and call-by-reference are two distinct methods for passing parameters to functions. In call-by-value, a copy of the actual data is passed, so any changes made inside the function do not affect the original variable. Call-by-reference, however, passes the original variable or its address, enabling the function to directly modify the caller's data, which is more efficient for large objects.

98. (A) Delegated Inheritance
C++ supports various types of inheritance to facilitate code reuse and hierarchical organization. Single inheritance involves one derived class from one base class, while multiple inheritance allows a class to inherit from several bases. Hierarchical inheritance involves multiple classes deriving from a single base. Delegated Inheritance is not a standard type of inheritance in C++, although delegation is a related design pattern used for composition.

99. (B) P-4, Q-1, R-3, S-2
Core C++ features provide unique benefits for software design. Encapsulation protects data through information hiding, while inheritance promotes code reuse and extension. Polymorphism offers flexibility by allowing different objects to share a common interface with dynamic behavior. Templates enable generic programming, allowing developers to write type-independent code that works efficiently across various data types without sacrificing performance or safety.

100. (C) (c)
Templates are the most suitable C++ feature for creating generic functions that operate on different data types using a single implementation. By defining the data type as a parameter, templates allow the compiler to generate specialized versions of a function for integers, floats, or doubles automatically. This approach significantly reduces code duplication

NODIA

Gouraud shading is a polygon rendering technique that produces smooth lighting by calculating color intensities at each vertex. These vertex intensities are then linearly interpolated across the surface of the polygon to determine the color of each interior pixel. This method eliminates the faceted appearance of flat shading, though it may not capture specular highlights as accurately as Phong shading.

47. (A) cerr : Unbuffered standard error stream
In C++, the cerr stream is specifically designed for reporting error messages to the standard error output. Unlike cout, which is buffered, cerr is unbuffered, ensuring that error information is displayed immediately as it occurs. This is critical for debugging and providing immediate feedback to the user or system logs during critical failures, where buffered output might be lost.

48. (B) Exception handling
Exception handling is the most effective mechanism in C++ for managing runtime errors, such as memory allocation failures. By using try and catch blocks, a programmer can attempt a risky operation and define a recovery path if it fails. This prevents the program from crashing unexpectedly and allows it to release resources or notify the user in a controlled manner.

49. (B) Both A and R are true but R is not the correct explanation of A.
The syntax analyzer, or parser, ensures that the sequence of tokens conforms to the language's grammar, building a hierarchical structure like an abstract syntax tree. While the creation of an AST is indeed a primary output of this stage, it is the process of grammatical verification itself that defines the role of the parser. Both aspects are fundamental to compilation.

50. (A) (d) only
In C++, virtual functions are designed specifically to be overridden in derived classes to support runtime polymorphism. Declaring a function as virtual allows a derived class to provide its own implementation, which is then dynamically called through a base class pointer. Other statements regarding destructors, multiple inheritance, templates, and constant member functions accurately reflect the standard features and rules of the language.

51. (C) The fill color and the boundary color.
The boundary-fill algorithm is used to fill a region by checking for a specific boundary color. Starting from an interior seed point, it recursively colors neighboring pixels until it encounters the defined boundary. This algorithm requires both the desired fill color and the color of the boundary itself to correctly identify the extent of the area that needs to be colored.

52. (D) Top to bottom, one scan line at a time.
A raster-scan display refreshes the screen by moving an electron beam horizontally across the screen, one scan line at a time, starting from the top. After each line, the beam returns to the left side and moves down to the next row. This process repeats rapidly to maintain a steady image, drawing the entire frame buffer contents onto the screen.

NODIA

53. (C) To determine which part of a line is inside a given rectangular window.
The Cohen-Sutherland algorithm is a line-clipping technique used to identify the segments of a line that fall within a defined rectangular viewing window. It assigns a four-bit region code to each line endpoint, allowing for quick rejection of lines completely outside the window and identifying lines that require intersection calculations with the window's boundaries for partial display.

54. (A) P-3, Q-1, R-2, S-4
Geometric transformations in 2D are represented by specific matrices. Scaling uses diagonal elements for the x and y factors. Rotation involves trigonometric functions for the specified angle. Reflection about the y-axis inverts the x-coordinate, and shearing shifts one coordinate based on the value of the other. These matrices are fundamental tools for manipulating objects within a computer graphics system.

55. (D) A class is a blueprint, while an object is an instance created from that blueprint.
In object-oriented programming, a class defines the structure, attributes, and methods that a particular type of entity will possess. It acts as a static template. An object is a dynamic, living instance of that class created at runtime. While the class exists in the source code, multiple

one pure virtual function, which mandates that derived classes provide a specific implementation. By preventing the creation of objects, abstract classes ensure that only concrete, fully defined subclasses can be used to generate functional program entities.

19. (B) (a), (b), (c), and (d) are the only correct items.
Bezier and B-Spline curves have distinct properties in computer graphics. A Bezier curve always starts and ends at its first and last control points, and its degree depends on the total number of points. B-Splines offer superior local control, meaning moving one point only affects a segment of the curve. These characteristics make B-Splines more flexible for complex modeling.

20. (D) virtual
The virtual keyword is used in a base class to declare a function that can be overridden by derived classes. This enables late binding, where the program determines which specific version of the function to execute at runtime based on the actual object type. This mechanism is essential for achieving polymorphism and creating flexible class hierarchies in object-oriented programming.

21. (C) Inheritance : A mechanism allowing a class to have multiple names.
Inheritance is a cornerstone of object-oriented programming that allows a new class to acquire the properties and behaviors of an existing class. This creates a hierarchical relationship, often described as an is-a relationship, which promotes code reusability and extensibility. By inheriting from a base class, a derived class can refine or expand functionality while maintaining a common interface.

22. (B) To display information in a visual form.
The primary function of a graphics monitor is to act as an output device that converts electronic signals from a computer into visible images. It provides a visual interface for users to interact with software, view graphical data, and observe the results of computations. Modern monitors use various technologies like LCD or OLED to produce high-quality color and detail.

23. (D) To map a selected part of a scene in world coordinates to device coordinates.
The Window-to-Viewport transformation is a critical step in the graphics pipeline that maps a specific region of world coordinates onto a designated area of the display device. The window defines what is being viewed in the scene, while the viewport determines where it appears on the screen. This process involves scaling and translation to ensure proper display.

24. (A) (iii), (i), (iv), (ii)
The translation process in a compiler typically follows a structured sequence. It begins with lexical analysis to group characters into tokens, followed by syntax analysis to check the program's grammatical structure. Semantic analysis then verifies the meaning and type consistency of the code. Finally, the code generation phase produces the executable machine instructions or intermediate code for the target system.

25. (C) Using an iterative method based on a decision parameter to select the next pixel.
The mid-point circle algorithm is an efficient technique for rasterizing circles by identifying the pixels that best approximate the circular path. It uses an incremental approach, calculating a decision parameter at each step to determine which adjacent pixel is closest to the ideal circle. This method relies primarily on integer arithmetic, making it computationally faster than trigonometric or square-root calculations.

26. (A) A data type
In the C language, tokens are the smallest individual units of a program that have meaning to the compiler. Keywords, identifiers, operators, constants, and string literals are all valid tokens. While data types like int or float are represented by keywords, the concept of a data type itself is an abstract classification rather than a specific lexical token in the source.

27. (D) A special member function that is automatically called when an object is created.
A constructor is a unique member function in C++ that shares the same name as its class. It is automatically triggered whenever an instance of the class is created. The primary purpose of a constructor is to initialize the object's data members and allocate necessary resources, ensuring that the object starts its lifetime in a valid and predictable state.

Downloaded from www.pdf.tube

65. (B) Runs on a web server to handle client requests and generate dynamic content.
A Java Servlet is a server-side component that extends the capabilities of a web server by dynamically processing client requests. It typically receives HTTP requests, interacts with databases or other services, and generates a response, often in HTML or JSON format. Servlets are a fundamental part of Java-based web development, providing a robust and portable way to build dynamic web applications.

66. (A) To simulate how light interacts with surfaces to determine their final color and brightness.
Illumination models are used in computer graphics to calculate the color and intensity of light at various points on an object's surface. These models consider factors such as light source position, surface orientation, and material properties like reflectivity and roughness. By simulating physical lighting effects, these models help create more realistic and visually appealing 3D images with depth and shading.

67. (A) Both Statement I and Statement II are true.
Flood-fill and boundary-fill are two common algorithms for filling regions in raster graphics. Flood-fill starts from a seed point and replaces a specific interior color with a new one. Boundary-fill also starts from a seed but continues coloring until it hits a specified boundary color. Both algorithms are recursive in nature and are used to fill closed areas in graphical images.

68. (D) Scaling
Scaling is a geometric transformation that alters the size of an object by multiplying its coordinates by specific factors. Uniform scaling changes the size equally in all directions, preserving the object's original proportions, while non-uniform scaling can stretch or shrink the object differently along different axes. This transformation is fundamental for resizing objects within a graphical scene to achieve the desired appearance.

69. (B) (iv), (i), (ii), (iii)
The 3D viewing pipeline follows a logical sequence to convert a 3D scene into a 2D image. It begins with modeling and viewing transformations to position objects and the camera. Next, a projection transformation maps the 3D coordinates onto a 2D plane. Clipping then removes parts outside the view volume, and finally, a viewport transformation scales the image for the display device.

NODIA

70. (A) Both A and R are true and R is the correct explanation of A.
Homogeneous coordinates allow all primary geometric transformations, including translation, to be represented as matrix multiplications. In standard Cartesian coordinates, translation is an additive process, which makes it inconsistent with multiplicative rotation and scaling. By adding an extra coordinate, translation becomes a multiplicative operation in the homogeneous system, enabling the concatenation of multiple transformations into a single, efficient composite matrix.

71. (D) A member function with the same name as the class, prefixed with a tilde (~).
A destructor is a special member function in C++ that is automatically invoked when an object's lifetime ends, such as when it goes out of scope or is explicitly deleted. It is identified by the tilde symbol followed by the class name. The primary role of a destructor is to perform cleanup tasks, such as releasing memory or closing file handles.

72. (A) Mapping 3D coordinates onto a 2D view plane.
Projection transformation is the process of converting three-dimensional coordinates of a scene into two-dimensional coordinates for display on a flat screen. This transformation can be perspective, where parallel lines converge, or parallel, where they remain parallel. It is a critical step in the 3D viewing pipeline that determines how objects in space are projected onto the viewer's image plane.

73. (C) Cohen-Sutherland shading
Polygon rendering involves techniques like Gouraud, Phong, and constant shading to determine the color and lighting of surfaces. These methods interpolate values across the polygon to create a smooth or shaded appearance. Cohen-Sutherland, however, is not a shading method but a line-clipping algorithm used to determine which parts of a line segment fall within a specific viewing window on the screen.

Homogeneous coordinates are a system used in computer graphics to represent N-dimensional points with N+1 coordinates. For a 2D point, this typically involves adding a third coordinate, usually set to one. This representation allows complex transformations like translation, rotation, and scaling to be expressed consistently as single matrix multiplications, greatly simplifying the overall computation in the graphics pipeline.

38. (B) They combine a sequence of transformations into a single matrix, improving efficiency.
Composite transformations allow multiple geometric operations, such as rotation, scaling, and translation, to be merged into a single mathematical matrix. By pre-multiplying these individual matrices, the system only needs to perform one matrix-vector multiplication for each point in the object. This significantly reduces the total number of calculations required, leading to faster rendering and improved graphical performance.

39. (B) (b) and (c) only
Scalar data types are fundamental, atomic units of data that cannot be broken down into simpler parts. Examples include integers and floating-point numbers, which represent single values. In contrast, composite types like arrays or structures are built from multiple scalar or other composite elements. Therefore, scalar types are characterized by their simplicity and lack of internal organizational structure or components.

40. (A) stdio.h
The stdio.h header file, standing for Standard Input Output, provides the declarations for the core file handling functions in the C language. This includes functions for opening, closing, reading from, and writing to files. By including this header, programmers gain access to the necessary tools for managing data persistence and interacting with the file system during program execution.

41. (C) To store an array of strings representing the command line arguments.
In a C program, the argv parameter is an array of character pointers, where each pointer points to a string passed from the command line. The first element usually contains the program's name, while subsequent elements contain the actual arguments provided by the user. This mechanism allows programs to receive input and configuration details directly upon startup from the terminal.

42. (A) The mechanisms for passing parameters to and returning values from functions.
Subprogram control encompasses the logic and protocols used to manage the execution of functions and procedures. This includes how data is passed into a function through parameters, how local variables are managed on the stack, and how results are returned to the calling code. It ensures that subprograms interact correctly, maintaining the integrity and flow of the overall program.

43. (D) Abstraction
Abstraction is a core principle of object-oriented programming that focuses on presenting the essential features of an entity while concealing the underlying implementation details. By providing a simplified interface, abstraction allows users to interact with complex objects without needing to understand their internal workings. This reduces complexity and improves software maintainability by decoupling the interface from the specific implementation.

44. (C) (a), (c), and (e) only
Web programming technologies serve different roles in the client-server architecture. Java Applets run on the client side, while Servlets execute on the server to handle requests. Scripting languages like JavaScript are generally interpreted by the browser for interactivity. Conversely, XML is used for data representation, not presentation, and DHTML is a general concept rather than a specific version of HTML.

45. (C) (x, -y)
Reflecting a 2D point across the x-axis involves keeping the horizontal coordinate the same while inverting the vertical coordinate. This operation creates a mirror image of the point on the opposite side of the x-axis. Mathematically, this is achieved by multiplying the y-coordinate by negative one, resulting in the new coordinates representing the reflected position in the Cartesian plane.

46. (D) Gouraud Shading

Click Here to Buy book on Amazon

83. (C) A program whose source code is split across several .cpp and .h files.
Multifile programs in C++ are organized by dividing the source code into multiple files, typically header files for declarations and implementation files for definitions. This approach improves code organization, facilitates modular development, and allows for separate compilation of different components. It makes large projects more manageable by enabling teams to work on different files simultaneously and reducing overall recompilation time.

84. (A) (a) and (d) only
In C++, a class can define multiple constructors with different parameter lists, a feature known as constructor overloading. This allows objects to be initialized in various ways. Every time an instance is created, the appropriate constructor is automatically triggered. However, a class can have only one destructor, which cannot take any arguments, as its role is simply to clean up the object.

85. (C) init()
The init() method is a fundamental part of a Java Applet's lifecycle and is executed exactly once when the applet is first loaded by the browser. Its purpose is to perform one-time initialization tasks, such as loading resources, setting initial parameters, or creating user interface components. Subsequent operations like starting and stopping the applet are handled by other lifecycle methods.

86. (B) It involves only integer arithmetic, making it computationally efficient.
The Mid-point Ellipse Algorithm is valued for its efficiency because it avoids complex floating-point calculations and trigonometric functions. By using an incremental approach based on an integer decision parameter, it identifies the best pixel approximations for an ellipse using only simple addition and multiplication. This makes it highly suitable for high-performance raster graphics systems where speed and resource usage are critical.

87. (D) Plotter
In computer graphics, input devices like mice, keyboards, and light pens are used to provide data and commands to the system. They allow users to interact with and manipulate graphical elements. A plotter, however, is an output device used to produce high-quality line drawings on paper. It receives data from the computer to create physical copies of images or designs.

88. (A) A coordinate system defined with the viewer's position as the origin.
The Viewing Coordinate Reference Frame is an intermediate coordinate system used in the 3D graphics pipeline. It centers the coordinate system on the viewer's position, with the axes oriented according to the direction the viewer is looking. This frame of reference simplifies the subsequent projection and clipping operations by aligning the scene relative to the observer's specific point of view.

89. (B) (b), (c), and (e) only
3D transformations are complex operations used to manipulate objects in space. Using 4x4 matrices with homogeneous coordinates allows translation to be represented as multiplication, just like rotation and scaling. Rotation can occur around any arbitrary axis, and shearing distorts shapes by shifting coordinates. However, translation requires more than a 3x3 matrix, and the order of transformations is critical as they are non-commutative.

90. (A) Members of the base class and its derived classes.
The protected access specifier in C++ provides a level of visibility that is intermediate between public and private. Members declared as protected are accessible within the base class and by any class derived from it. However, they remain hidden from other parts of the program, preserving encapsulation while still allowing subclasses to utilize and extend the base class's internal state and behavior.

91. (D) Cohen-Sutherland algorithm
The Cohen-Sutherland algorithm is a classic line-clipping technique that uses four-bit region codes, or outcodes, to categorize line endpoints relative to a rectangular window. Each bit represents whether a point is above, below, to the left, or to the right of the window. By comparing these outcodes, the algorithm can quickly determine if a line is completely visible, invisible, or needs clipping.

NODIA

objects can be created from it, each holding its own unique data.

56. (A) Environmental
Programming paradigms represent the fundamental models of computation, such as imperative, declarative, and object-oriented styles. Imperative programming focuses on state changes and commands, while declarative programming emphasizes what the program should accomplish. Object-oriented programming organizes logic around data objects. Environmental is not a recognized programming paradigm, although it may refer to the context in which a program executes or is developed.

57. (C) Domain
The domain of a data type defines the complete set of all possible values that a variable of that type can hold. For instance, the domain of a boolean type consists only of true and false. Understanding the domain is essential for ensuring type safety, as it dictates the valid inputs and outputs for operations performed on data of that specific type.

58. (A) To include the content of another file into the source file.
The #include directive is a preprocessor command used in C and C++ to insert the entire contents of a specified file into the current source file at the location of the directive. This is commonly used to include header files that contain function declarations, macro definitions, and class structures, promoting code modularity and allowing reuse of standard library components.

59. (A) (a) only
Valid identifiers in the C language must follow specific naming rules. They must start with a letter or an underscore and can be followed by letters, digits, or underscores. Identifiers cannot begin with a digit, contain special characters like hyphens, or be reserved keywords. Consequently, an identifier starting with an underscore is valid, while those starting with digits or using hyphens are not.

60. (B) The behavior is undefined, often leading to a resource leak.
In C++, when a base class pointer points to a derived class object, deleting that pointer requires a virtual destructor in the base class. If the destructor is not virtual, only the base class destructor is called, leaving the derived portion of the object uncleaned. This leads to undefined behavior and potential memory leaks, as resources allocated by the derived class are never released.

61. (C) The process of creating an instance (an object) of a class.
Instantiation is the action of producing a concrete object based on the definition provided by a class. When a class is instantiated, the system allocates memory for the object's attributes and invokes its constructor to initialize its state. This process transforms an abstract class blueprint into a functional entity that can participate in the execution and logic of the program.

62. (D) An active edge table (AET) sorted by x-coordinates.
The scan-line polygon fill algorithm relies on an active edge table to efficiently track the edges that intersect the current scan line. This table stores information about the x-intersections and slopes, and it is sorted by x-coordinates to facilitate filling the spans between pairs of intersections. This data structure ensures that the algorithm processes the polygon systematically and correctly.

63. (A) It provides local control, where moving one control point only affects a part of the curve.
Local control is a key advantage of B-spline curves, distinguishing them from Bezier curves. In a B-spline, each control point only influences a specific segment of the curve defined by the basis functions. This allows designers to make precise adjustments to one part of a curve without inadvertently changing its entire shape, providing greater flexibility and control in geometric modeling.

64. (A) (a), (b), and (d) only
A reference in C++ acts as an alternative name for an existing variable and must be initialized upon declaration. Once established, a reference cannot be redirected to another variable and occupies the same memory address as the original. Unlike pointers, references cannot be null, ensuring that they always refer to a valid object, which provides a safer and more intuitive syntax.

Downloaded from www.pdf.tube

while maintaining strict type safety during compilation.

101. (C) Raster targets the whole screen; random draws only where needed.
Raster-scan displays refresh the entire screen by scanning an electron beam line by line, regardless of the content. In contrast, random-scan displays direct the beam only to the specific coordinates where the image components, such as lines and characters, are to be drawn. This fundamental difference makes random-scan displays more efficient for line drawings but less capable of displaying realistic, shaded images.

102. (B) Viewing transformation
Viewing transformation is the process of mapping a selected portion of a world-coordinate scene onto a specific area of a display device. This involves defining a window in the world and a viewport on the screen. The transformation ensures that the scene is correctly scaled and positioned within the viewport, allowing users to view a specific part of a larger graphical world.

103. (D) To allow multiple functions with the same name but different parameter lists to exist.
Function overloading is a feature in C++ that allows several functions to share the same name as long as their parameter lists are unique in terms of number or types of arguments. This enables programmers to use a single name for operations that are conceptually similar but work with different inputs, improving code readability and providing a more intuitive interface.

104. (B) (a), (b) and (c) only
C++ provides specialized stream classes for file operations, with ofstream for writing to files, ifstream for reading from them, and fstream supporting both. These classes are typically part of the <fstream> header, which must be included for file I/O. While <iostream> is used for standard console input and output, the specific file stream classes are essential for managing persistent data on disk.

105. (C) Interface to access and modify document structure and content.
The Document Object Model (DOM) is a programming interface for web documents that represents the structure of an HTML or XML document as a tree of objects. Scripting languages like JavaScript use the DOM to dynamically access, add, change, or delete elements and content within a web page. This allows developers to create highly interactive and responsive user experiences in real-time.

106. (A) Orthographic Projection
Orthographic projection is a type of parallel projection where parallel lines in 3D space remain parallel in the projected 2D image. This is achieved by projecting all points along lines that are perpendicular to the view plane. This method is commonly used in engineering and architectural drawings because it preserves accurate dimensions and relative proportions, although it lacks the depth cues of perspective.

107. (D) A mathematical representation of a curve or surface defined by a set of control points.
Spline representation is a mathematical technique used in computer graphics to define smooth, continuous curves and surfaces using a set of control points. These points influence the shape of the spline without the curve necessarily passing through all of them. Splines provide a flexible and precise way to model complex shapes in design, animation, and various engineering applications, ensuring high-quality geometric representation.

108. (B) Phong model interpolates surface normals across a polygon and calculates lighting for each pixel, producing better specular highlights.
The Phong illumination model offers a more realistic representation of lighting than Gouraud shading by interpolating surface normal vectors across the polygon instead of color intensities. By calculating the lighting model at every pixel using these interpolated normals, the Phong model accurately captures specular highlights and subtle reflections, resulting in a smoother and more detailed appearance of curved surfaces in 3D graphics.

NODIA

109. (A) (a), (b), and (c) only
A fundamental programming environment consists of the essential tools needed to create, translate, and test software. This includes a text editor for writing source code, a compiler or interpreter to convert it into an executable format, and a debugger to identify and resolve

the web browser. By restricting access to sensitive resources like the local file system and private networks, the sandbox prevents potentially malicious applets from causing harm to the user's computer. This model was essential for allowing untrusted code to run securely as part of a web page.

165. (C) Eye Coordinate System
The Viewing Coordinate Reference Frame is frequently referred to as the Eye Coordinate System because it is centered on the viewer's position. In this system, the origin represents the observer's eye, and the axes are aligned with their line of sight and orientation. This frame of reference is a critical intermediate step in the graphics pipeline, simplifying the processes of projection and clipping against the viewer's field.

166. (B) Phong Shading
Phong shading is an advanced polygon rendering technique that produces high-quality lighting effects by interpolating surface normal vectors across each polygon. By calculating the illumination model for every pixel using these interpolated normals, the method accurately captures specular highlights and subtle reflections. This results in a much smoother and more realistic appearance for curved surfaces in 3D graphics compared to Gouraud or flat shading.

167. (D) auto : Preprocessor
In the C language, tokens are categorized into several types, including identifiers, operators, and constants. While printf is an identifier for a library function, + is an operator and 12.5 is a floating-point constant. The auto keyword, however, is a storage class specifier, not a preprocessor directive. Preprocessor directives, such as #include or #define, are handled before the actual compilation process begins.

168. (A) (iii), (ii), (i)
Variable lifetimes in C are determined by their storage class. Local automatic variables have the shortest lifetime, existing only while the function they are in is executing. Static variables within a function last for the entire duration of the program's execution, but their scope is restricted. Global variables also exist for the program's full lifetime and are accessible from anywhere, representing the most persistent data.

169. (A) P-4, Q-3, R-2, S-1
C++ uses specific constructs to manage the lifecycle and visibility of objects and their members. A constructor initializes an object upon creation, while a destructor handles cleanup during its destruction. The this pointer provides a way to refer to the current instance within member functions. Additionally, a friend function is a non-member that is granted access to a class's private and protected data.

170. (A) (a) only
Templates in C++ are a powerful mechanism for achieving compile-time, or static, polymorphism rather than runtime polymorphism. They allow developers to write generic code that can work with various data types, with the compiler generating the specialized versions as needed. This significantly reduces code duplication and improves maintainability while maintaining high performance, which is a key goal of generic programming in C++.

171. (D) Pixel
A pixel, or picture element, is the most fundamental building block of a raster-scan display's image. The screen is composed of a grid of these tiny individual points, each of which can be assigned a specific color and intensity. By combining thousands or millions of pixels, the display can represent complex and detailed images, which are stored in the frame buffer and refreshed regularly.

172. (B) The degree to which x-coordinates are shifted based on the y-coordinate.
In a 2D shear transformation parallel to the x-axis, the shx parameter determines the amount of horizontal shift applied to each point. This shift is proportional to the point's vertical coordinate, meaning points higher up are shifted more than those near the axis. This operation distorts the shape of the object, turning a rectangle into a slanted parallelogram while keeping its area and height constant.

173. (C) To identify and discard parts of a picture that are outside the viewing window or volume.
Clipping is a critical operation in computer graphics that ensures only the relevant parts of a scene are rendered for the user. It identifies the portions of geometric primitives, such as lines

Click Here to Buy book on Amazon

pointer, operator overloading for indexing, and the role of friend functions.

119. (C) A is true but R is false.
Encapsulation is the fundamental OOP principle that bundles data and methods into a single class, providing data protection and information hiding. However, the primary goal of encapsulation is to manage complexity and ensure data integrity, not specifically to facilitate code reuse through inheritance. While encapsulation and inheritance are related, each serves a distinct purpose within the overall object-oriented model.

120. (A) The number of command line arguments plus one (for the program name).
In a C program, the argc parameter represents the total count of arguments passed to the program from the command line, including the name of the executable itself. This value is used to determine how many entries are in the argv array, allowing the program to process the provided input parameters correctly. It is a standard mechanism for command-line interaction.

121. (C) Logic Programming
Logic programming is a paradigm based on formal logic where programs are composed of facts and rules. A logic-based system, such as Prolog, uses these definitions to answer queries by performing automated reasoning. This approach is distinct from imperative or object-oriented programming, as it emphasizes the relationships between data points rather than the specific sequence of steps to achieve a result.

122. (A) struct
A struct, or structure, is a prime example of a composite data type because it allows the grouping of multiple variables, potentially of different types, into a single unit. This contrasts with scalar types like char or float, which represent single, atomic values. Composite types like structures and arrays are essential for building complex data models and organizing information logically within a program.

123. (D) Subprogram Control
Subprogram control refers to the mechanisms used to define, invoke, and manage functions or procedures within a program. This includes the definition of parameters, the execution of the function's logic, and the returning of a value to the caller. By using subprograms, developers can break complex tasks into smaller, manageable, and reusable components, which is a fundamental practice in structured programming.

124. (A) P-3, Q-4, R-1, S-2
Object-oriented programming is built upon several core concepts. A class serves as a blueprint or template, while an object is a specific, tangible instance created from that blueprint. Methods are the functions associated with a class that define its behavior, and attributes are the data members or properties that describe the object's state. Together, these elements form the basis of the OOP model.

125. (C) Function Overriding
Function overriding is a feature in object-oriented programming that allows a derived class to provide a specific implementation for a method that is already defined in its base class. This is a key mechanism for achieving runtime polymorphism, as it enables the program to call the most specialized version of a method based on the actual type of the object at execution time.

126. (D) HTML tags are predefined, while XML allows you to define your own tags.
The most significant difference between HTML and XML lies in their tag sets. HTML uses a fixed set of predefined tags designed specifically for displaying content in a web browser. XML, on the other hand, is a meta-language that allows developers to create their own custom tags to describe the structure and meaning of the data they are storing or transporting.

127. (A) Color
Graphics primitives are the basic geometric building blocks used to create complex images, such as points, lines, and polygons. These entities have geometric properties like position and size. Color, however, is an attribute applied to these primitives rather than a primitive itself. It defines the visual appearance and intensity of the shapes drawn on the screen within a graphics system.

128. (B) To determine the visible portions of a line that are inside a convex clipping window using parametric equations.

Downloaded from www.pdf.tube

While both algorithms are used for region filling, they use different criteria to define the fill area. The boundary-fill algorithm continues coloring until it encounters a specific boundary color. In contrast, the flood-fill algorithm is used to replace an existing interior color with a new one throughout a connected region. This distinction makes each algorithm suitable for different types of graphical filling tasks.

183. (D) (a), (b), (c) and (d)
Geometric transformations are fundamental tools in graphics with well-defined mathematical properties. Rigid body transformations maintain the distance and angles of an object. Concatenating transformation matrices into a single composite matrix is highly efficient, although the order of application is critical because matrix multiplication is non-commutative. Uniform scaling is essential for changing an object's size while perfectly preserving its original shape and proportions.

184. (A) Both A and R are true and R is the correct explanation of A.
The const keyword is a valuable feature in C++ that allows the declaration of immutable variables. By using const, programmers can prevent accidental changes to data that should remain constant, leading to safer and more predictable code. Additionally, the compiler can use this information to perform various optimizations, such as constant folding, which can improve the performance and efficiency of the final program.

185. (C) It defines the region of 3D space that is visible in the final image.
The view frustum, or viewing funnel, is a three-dimensional volume that represents the field of vision for a camera or observer in a 3D scene. Only the objects or parts of objects that fall within this volume are projected onto the 2D view plane and rendered. The frustum is defined by the near and far clipping planes and the boundaries of the viewing window.

186. (A) Polygon Spline
Spline representations in computer graphics include various mathematical models for defining smooth curves, such as Bezier, B-spline, and cardinal splines. These splines use control points to influence the curve's shape. While polygons are used to represent surfaces as meshes of flat facets, there is no standard Polygon Spline representation, as splines are fundamentally defined by parametric or algebraic functions rather than simple line segments.

187. (D) Small set of primitives combined to build complex structures.
Orthogonality in programming language design refers to the principle that a relatively small number of basic constructs can be combined in a limited number of ways to build any complex program structure. A highly orthogonal language is easier to learn and use because the behavior of combined elements is consistent and predictable, with few exceptions or special cases to remember.

188. (B) Pass a pointer to the structure.
Passing a pointer to a large structure is the most efficient method because only the memory address of the structure is passed to the function, rather than a full copy of all its data. This significantly reduces memory overhead and improves performance, especially for complex or deeply nested structures. It also allows the function to directly modify the original structure if needed.

189. (B) P-2, Q-1, R-4, S-3
C++ uses various operators for accessing members and memory. The dot operator is for object variables, while the arrow operator is used when dealing with pointers to objects. The scope resolution operator resolves the context for identifiers in different namespaces or classes. The ampersand operator is used to retrieve the memory address of any variable, which is essential for pointer initialization and reference handling.

190. (A) (d) only
In C++ exception handling, the try block identifies risky code, and the throw statement raises an exception. If caught, execution transfers to the catch block. After the catch block finishes, the program continues from the statement following the catch block, not from the point where the exception was thrown. The throw statement effectively jumps out of the normal execution flow to find a handler.

137. (C) Pointer to a constant string literal.
In the C language, declaring a character pointer and assigning it a string literal creates a pointer that stores the memory address of the first character in that string. The string itself is typically stored in a read-only section of memory, making it a constant. This approach is commonly used for handling text strings efficiently by passing addresses instead of entire arrays.

138. (A) The ability to process objects of different types through a single, uniform interface.
Polymorphism is a core principle of object-oriented programming that allows objects of different classes to be treated as if they belong to a common base class. This is achieved through mechanisms like function overriding and virtual functions, which enable the program to invoke the correct specific behavior at runtime, providing a flexible and consistent way to interact with varied object types.

139. (C) A is true but R is false.
The Sutherland-Hodgman algorithm is a standard technique for clipping a polygon against each edge of a convex window sequentially. While it is efficient and widely used, it has a notable limitation: it may produce a single, connected set of vertices even when the clipping process should logically result in multiple separate, disjoint polygons. This can lead to extraneous lines connecting the separated regions.

140. (B) (a) and (d) only
Virtual functions in C++ are essential for implementing runtime polymorphism and dynamic binding. By using the virtual keyword in the base class, the programmer ensures that when a function is called through a base pointer, the version corresponding to the actual object type is executed. This is distinct from static polymorphism, which is achieved through function overloading and resolved at compile time.

141. (C) Handling exceptions and preventing program termination from runtime errors.
A try-catch block is the standard C++ mechanism for exception handling. The try block encloses code that might generate a runtime error, while the catch block provides the logic to handle the exception if it occurs. This structure allows the program to recover gracefully from errors instead of terminating abruptly, which is crucial for building reliable and user-friendly software applications.

142. (D) Scaling and Translation
The Window-to-Viewport transformation is a process that maps coordinates from a viewing window in the world to a specific viewport on a display device. This operation primarily involves scaling to adjust the size of the objects to fit the viewport and translation to position the mapped content correctly on the screen. Together, these steps ensure the selected scene is properly displayed.

143. (A) (a), (b), and (c) only
3D objects can be represented using various techniques, each with its own advantages. Polygon surfaces use meshes for efficient rendering, while quadric surfaces provide mathematical precision for spheres and cylinders. Spline surfaces, such as Bezier and B-spline surfaces, are essential for modeling smooth, curved shapes with great flexibility and local control, making them indispensable in design and engineering applications.

144. (B) P-3, Q-2, R-1, S-4
Graphics algorithms serve diverse functions in rendering and image processing. Bresenham's algorithm provides efficient line drawing using integer arithmetic, while Cohen-Sutherland handles line clipping. The boundary-fill algorithm is used to fill regions starting from an interior point. Gouraud shading produces smooth surface lighting by interpolating color intensities across polygons. These techniques are fundamental for creating high-quality images in computer graphics.

145. (B) The source code is checked to ensure it conforms to the grammatical rules of the language.
Syntax analysis, also known as parsing, is a crucial stage of compilation where the sequence of tokens produced by the lexical analyzer is checked against the formal grammar of the programming language. This phase ensures that the program is structurally correct and follows all grammatical rules. The result is typically a

Click Here to Buy book on Amazon

hierarchical structure, such as a parse tree or an abstract syntax tree.

146. (A) float temperature
A scalar data type is a fundamental, atomic unit that holds a single value, such as an individual number or character. A floating-point variable like temperature is a classic example. In contrast, composite types like structures, arrays, and unions are built by combining multiple scalar or other composite elements into a single organized unit. Scalar types are the basic building blocks of all data.

147. (C) A pointer p to an integer.
In the C language, the asterisk symbol in a variable declaration specifies that the variable is a pointer. Therefore, declaring a variable as an integer pointer means it will store the memory address of an integer variable rather than an integer value itself. Pointers are essential for dynamic memory allocation, efficient array handling, and implementing complex data structures like linked lists.

148. (D) Polymorphism
The ability of different classes to provide their own specific implementations of a method defined in a base class is a prime example of polymorphism. By having each shape class implement its own draw() function, the program can interact with any shape through a common interface. The correct specific version of the method is then determined and executed at runtime based on the actual object type.

149. (A) (d) only
In C++, access specifiers like public, private, and protected are used to control the visibility and accessibility of class members. Public members are accessible from anywhere, private members only within the class, and protected members by derived classes. The virtual keyword, however, is not an access specifier but a modifier used to enable dynamic binding and runtime polymorphism for member functions.

150. (B) (b)
The <form> tag is the primary container for creating interactive forms in HTML. It encloses all the individual input elements, such as text boxes, buttons, and labels, and defines where and how the data should be sent when the form is submitted. This tag is essential for facilitating user input and data collection on web pages, providing a structured way to handle information.

151. (D) A vector-refresh display
A vector-refresh display is a classic example of a random-scan system, where the electron beam is directed to draw lines directly on the screen rather than scanning the entire surface line by line. This technology was historically used for high-quality line drawings in engineering and scientific applications. Because it traces geometric primitives directly, it produces very smooth lines without the pixelation of raster-scan systems.

152. (A) [[0, 1], [1, 0]]
Reflecting a 2D point across the line y = x involves swapping its x and y coordinates. This transformation is mathematically represented by a specific matrix where the off-diagonal elements are one and the diagonal elements are zero. When this matrix is multiplied by a point's coordinates, the horizontal and vertical positions are interchanged, creating a mirror image across the diagonal line.

153. (C) A view point (eye), a look-at point, and a view-up vector.
The Viewing Coordinate Reference Frame is defined by three key parameters: the position of the viewer (the eye point), the direction they are looking (the look-at point), and the orientation of their head (the view-up vector). These parameters together establish a coordinate system centered on the observer, which is essential for correctly projecting and rendering the 3D scene from their perspective.

154. (A) (a), (b), (c), and (e) only
3D projections serve different purposes in computer graphics. Perspective projection provides a realistic view by making distant objects appear smaller, with parallel lines converging at vanishing points. Orthographic projection is a parallel projection where parallel lines remain parallel, often used for technical drawings. For orthographic projections, the center of projection is effectively at infinity. Oblique projection is also a type of parallel projection, not perspective.

155. (D) Code Generation

NODIA

Code generation is the final stage of the compilation process, where the compiler transforms the analyzed and optimized representation of the source program into the actual target machine code or intermediate language. This stage involves mapping program constructs to specific processor instructions and managing hardware resources like registers and memory addresses, resulting in the final executable file that can be run on the target system.

156. (A) Both Statement I and Statement II are true.
Data types are fundamentally defined by the set of values they represent and the operations allowed on those values. This ensures that data is used consistently and safely within a program. Furthermore, composite data types are created by combining simpler types, such as scalars or other composites, into more complex structures like arrays or records, allowing for more advanced data organization.

157. (C) 8, 7
In the C language, the sizeof operator returns the total memory occupied by a variable, including the hidden null terminator at the end of a string array. For the string PROGRAM, which has seven characters, the array size is eight. Conversely, the strlen function counts only the actual characters before the null terminator, returning seven. This distinction is critical for memory management.

158. (D) Hierarchical Inheritance
Hierarchical inheritance occurs when multiple classes are derived from a single base class. In this model, each derived class inherits the common properties and behaviors of the parent class while adding its own specific features. This is a common design pattern in object-oriented programming for creating specialized versions of a general concept, such as different types of vehicles sharing a common base class.

159. (B) (i), (ii), (iii), (iv)
C++ features are often built upon one another in a logical hierarchy. A class must be defined before an object can be instantiated from it. Inheritance then allows a new class to derive from an existing one, and virtual functions are used within an inheritance hierarchy to enable dynamic behavior. Understanding this sequence is essential for mastering the structural and functional aspects of object-oriented programming.

160. (A) Both A and R are true and R is the correct explanation of A.
The try...catch structure provides a robust framework for managing runtime exceptions in C++. When an error is detected within the try block, the program throws an exception, which immediately transfers control to a corresponding catch block. This structured approach allows developers to handle errors locally or propagate them up the call stack, ensuring that the program can recover or terminate gracefully.

161. (B) Dynamic Hypertext Markup Language
DHTML stands for Dynamic Hypertext Markup Language, which is a collection of technologies used to create interactive and animated websites. It typically combines standard HTML for content structure, CSS for visual styling, and JavaScript for client-side behavior. Together, these tools allow a web page's content and appearance to be modified dynamically in response to user actions without needing to reload the entire page.

162. (A) The implicit form of the circle equation, $f(x, y) = x^2 + y^2 - r^2 = 0$.
The mid-point circle algorithm utilizes the implicit equation of a circle to determine which pixels should be illuminated on a raster display. By evaluating this function at the midpoint between potential pixel choices, the algorithm derives a decision parameter. The sign of this parameter indicates whether the midpoint is inside or outside the ideal circle, guiding the selection of the closest pixel approximation.

163. (A) (a), (b), and (d) only
Bezier curves are defined by a set of control points, with the curve's degree being one less than the number of points. They are tangent to the first and last segments of the control polygon and exhibit global control, meaning moving any point affects the whole curve. However, the curve only passes through the first and last points, not necessarily all intermediate ones.

164. (A) Both A and R are true and R is the correct explanation of A.
The Java applet security model uses a sandbox to provide a safe execution environment within

Downloaded from www.pdf.tube

errors. While other tools like version control and specific hardware are important, these three components are the core requirements.

110. (C) A is true but R is false.
The sizeof operator is a valuable tool in C for determining the memory footprint of a variable or data type in bytes. However, the exact size of primitive types like integers is not universally fixed; it can vary depending on the computer architecture and the specific compiler being used. Therefore, assuming that the size of an integer is constant across all systems is incorrect.

111. (A) At run time
Dynamic binding is a mechanism where the association between an identifier and its properties, such as a function call and its implementation, is resolved during the program's execution. This allows for more flexible and adaptable code, as the exact behavior can be determined based on the state of the program or the specific objects being processed at runtime, which is essential for polymorphism.

112. (B) All members of a structure share the same memory location.
In C, a structure is a composite data type that allows grouping variables of different types under a single name. Each member of a structure is allocated its own distinct memory space, allowing them all to be accessed and modified independently. This is the primary difference from a union, where all members share the same memory location, meaning only one member is active at once.

113. (B) In the class destructor.
The class destructor is the most reliable place to put code for closing resources like network connections or file handles. Since the destructor is automatically called whenever an object's lifetime ends, it ensures that the connection is properly closed regardless of how the object is destroyed. This practice, known as Resource Acquisition Is Initialization (RAII), prevents resource leaks and improves software robustness.

114. (A) (ii), (iii), (i), (iv)
The 2D viewing pipeline systematically transforms a scene for display. It begins by constructing the world-coordinate scene using modeling transformations. Next, it converts these world coordinates to a viewing coordinate system centered on the observer. Clipping then removes any parts outside the defined viewing window. Finally, the remaining portion is mapped to the display device's normalized viewport coordinates for final rendering.

115. (C) [[cosθ, -sinθ], [sinθ, cosθ]]
A 2D rotation around the origin is mathematically represented by a specific transformation matrix. This matrix uses cosine and sine functions of the rotation angle to recalculate the coordinates of each point. By multiplying the original coordinates by this matrix, the object is rotated counter-clockwise around the origin. This matrix is a fundamental component of the geometric transformation process in computer graphics.

116. (C) Statement I is false, but Statement II is true.
The degree of a Bezier curve is determined by the number of its control points minus one. A primary advantage of B-spline curves over Bezier curves is the property of local control, where moving a single control point only affects a specific segment of the curve. This makes B-splines much more flexible and easier to manipulate for complex geometric modeling and design tasks.

117. (C) Specular reflection
Specular reflection is a key component of the Phong illumination model, representing the bright, mirror-like highlights seen on shiny surfaces. This component depends on the viewer's position and the orientation of the surface normal. By incorporating specular reflection along with ambient and diffuse lighting, the Phong model creates a more realistic and visually appealing simulation of how light interacts with smooth, reflective materials.

118. (A) (b) and (e) only
In C++, several statements about class members and operators are factually incorrect. Static member functions can only access static data members directly, not non-static ones, as they lack a this pointer. Furthermore, new and delete are operators, not functions, and provide built-in mechanisms for dynamic memory management. Other statements correctly describe the this

NODIA

and polygons, that fall outside the boundaries of a defined viewing window or three-dimensional volume. By discarding these invisible parts, the system avoids unnecessary calculations and correctly presents the visible world.

174. (A) (a), (b), and (c) only
Effective programming language syntax is designed for high readability and writability while remaining unambiguous to ensure consistent interpretation by compilers. Orthogonality is another key principle, where a small set of basic constructs can be combined in predictable ways to build complex logic. While efficiency is important, mimicking machine code is generally not a goal of high-level language syntax, as it reduces abstraction and portability.

175. (B) (a), (b), d, (e) only
Elementary data types form the foundation of data representation in programming. Scalar types are simple, atomic units, while composite types are built from other types. Every data type is characterized by its name, value domain, and the operations it supports. Furthermore, the type dictates how information is structured and stored in memory, though a struct in C is a composite type, not a scalar.

176. (B) To indicate that the pointer does not point to any valid memory location.
The NULL pointer is a special constant in C used to signify that a pointer variable is not currently directed to any valid or meaningful memory address. It is a vital tool for initializing pointers and checking whether they have been assigned a proper location before use. This practice helps prevent serious runtime errors, such as accessing unallocated memory or dereferencing invalid pointers.

177. (B) Multilevel inheritance
Multilevel inheritance is a type of class hierarchy where a class is derived from another derived class. This creates a chain of inheritance, such as class D inheriting from class B, which in turn inherits from class A. In this model, the most derived class gains the features and properties of all its ancestors in the chain, facilitating a hierarchical organization and reuse of code.

178. (A) To specify a class's scope for a member or to access global variables (Scope Resolution Operator).
The scope resolution operator is used in C++ to explicitly define the context for an identifier. It allows programmers to access a global variable even if a local variable with the same name exists, or to define a member function outside of its class declaration. This operator is essential for managing namespaces and resolving potential naming conflicts in complex programs with multiple scopes.

179. (A) (a), (b), and (d)
Runtime dynamic behavior is achieved through a combination of several core object-oriented features. Inheritance provides the structural hierarchy, while virtual functions enable dynamic binding by allowing methods to be overridden. Polymorphism allows the program to interact with different objects through a common interface, with the correct behavior being determined at execution time. Together, these mechanisms provide the flexibility required for dynamic systems.

180. (A) P-3, Q-4, R-1, S-2
Programming languages are often associated with the specific paradigms they primarily support. C is a classic procedural language, while C++ is known for its robust object-oriented features. LISP is a foundational functional language that emphasizes list processing and recursion. Prolog is the most prominent logic programming language, using a declarative approach based on facts and rules to perform automated reasoning and problem-solving.

181. (C) Java Servlet
A Java Servlet is a robust, server-side technology used to extend the functionality of a web server. It provides a platform-independent way to handle client requests and generate dynamic responses, such as custom HTML pages or data streams. Servlets are a key part of Java-based enterprise web development, offering a scalable and efficient alternative to older server-side scripting methods for building complex web applications.

182. (B) Flood-fill checks for an interior color to replace, while boundary-fill checks for a boundary color.

Click Here to Buy book on Amazon

The Liang-Barsky algorithm is an efficient line-clipping technique that uses parametric equations to represent a line segment. It identifies the portions of the line that fall within a convex clipping window by solving for the parameter values at the boundaries. This approach is often faster than the Cohen-Sutherland algorithm because it reduces the number of intersection calculations required for clipping.

129. (B) (b) and (d) only
Inheritance in C++ involves specific access rules for members of the base class. Public inheritance establishes an is-a relationship, making public members accessible. Protected members are accessible to derived classes but not to outside code, while private members remain hidden even from derived classes. Additionally, constructors and destructors are not inherited but are called in a specific sequence during object creation and destruction.

130. (A) Event Handling
Event handling is the programming mechanism used to respond to user actions, such as mouse clicks, key presses, or menu selections. In a graphical user interface, these actions generate events that the program must capture and process to perform the intended tasks. This concept is fundamental for creating interactive software where the program's flow is determined by user-initiated interactions.

131. (C) It draws lines and characters directly, leading to smooth lines.
Random-scan displays, also known as vector displays, operate by directing the electron beam to trace lines and characters directly between specified coordinates. This results in very smooth lines without the jagged edges often seen in raster-scan systems. Because the beam only visits the parts of the screen where an image is present, these displays are highly efficient for complex line-based graphical representations.

132. (D) Translation
Translation is a basic geometric transformation that moves an object from one location to another in a straight line. This is achieved by adding specific displacement values to the coordinates of every point in the object. Unlike scaling or rotation, translation does not change the object's size, shape, or orientation; it only modifies its position within the coordinate system or scene.

133. (C) To define the sequence of operations for creating a 2D image from a 3D scene description.
The viewing pipeline is a series of stages that a graphics system follows to transform a three-dimensional scene into a two-dimensional image for display. This includes modeling, viewing, projection, clipping, and viewport transformations. By systematically processing geometric data through these stages, the system can correctly represent the 3D world from a specific viewpoint onto a flat screen or display device.

134. (A) (a) and (d) only
Polygon clipping involves specialized algorithms like Sutherland-Hodgman and Weiler-Atherton. Sutherland-Hodgman is designed for convex windows and processes the entire polygon against each window edge. Weiler-Atherton is more advanced, capable of correctly clipping both convex and concave polygons and handling disjoint regions. Polygon clipping is generally more complex than line clipping because it must maintain the integrity of closed shapes and potential new boundaries.

135. (D) A collection of tools and software to aid in the development of programs.
A programming environment is an integrated set of software tools that assist developers in the entire process of writing, testing, and debugging code. This typically includes a text editor, a compiler or interpreter, a debugger, and often version control and build management systems. A well-designed environment improves productivity by providing a streamlined and consistent workflow for software creation and maintenance.

136. (A) Its type, which determines the operations that can be performed on it.
In the context of data types, an object's type is its most fundamental property, as it defines the set of values it can take and the specific operations that can be applied to it. For example, an integer type allows arithmetic operations, while a string type might allow concatenation. The type ensures data integrity and determines how the system interprets and manipulates the information.

NODIA

191. (A) Plasma Panel
A plasma panel display is a type of flat-panel display that uses small cells containing ionized gases to create an image. These gases respond to electric fields by emitting light. This technology allowed for larger and thinner displays compared to traditional cathode ray tubes. By managing the charge distribution across the grid of capsules, the system can produce high-quality color images with excellent contrast and wide viewing angles.

192. (B) Cyrus-Beck
The Cyrus-Beck algorithm is a more general and powerful line-clipping technique than Cohen-Sutherland, as it can clip a line segment against any convex polygon, not just a rectangular window. It uses a parametric representation of the line and the normal vectors of the clipping edges to determine the intersections. This makes it a versatile tool for more complex clipping scenarios in computer graphics and geometric modeling.

193. (B) To remove objects or parts of objects that are outside the view frustum.
Clipping in the 3D viewing pipeline is the process of identifying and discarding geometric data that falls outside the defined view frustum. By removing these invisible parts, the system avoids the computational expense of processing and rendering objects that would not appear in the final 2D image. This optimization is essential for maintaining high performance and ensuring the correct visual representation of the scene.

194. (D) (a), (b), (c) and (d)
Valid variable declarations in C cover a range of types and initializations. These include simple integer assignments, constant floating-point values for mathematical precision, character arrays for storing text, and structures for grouping related data. Each of these declarations follows the standard syntax and rules for defining data within a C program, providing the necessary foundations for managing varied types of information.

195. (C) To provide a mathematical or abstract model for the behavior of a computational system, often used in compiler design.
A formal transition model provides a rigorous mathematical framework for describing how a system changes states in response to inputs. In the context of language design and compilation, these models, such as finite automata or pushdown automata, are essential for formally specifying the lexical and syntactic rules of a language. They ensure that compilers can consistently and accurately translate source code into executable programs.

196. (C) Late Binding
Late binding, also known as dynamic binding, is a programming mechanism where the specific method to be executed is determined at runtime based on the actual type of the object. This is a fundamental part of polymorphism in object-oriented programming, allowing a common interface to support diverse behaviors. It contrasts with early binding, where the association is fixed by the compiler during the translation phase.

197. (C) Program that handles directives before the actual compilation.
The C preprocessor is a specialized tool that processes the source code before it reaches the compiler. It handles directives like #include for file insertion, #define for macro substitution, and #ifdef for conditional compilation. By performing these text-based transformations, the preprocessor prepares the final source file for the actual compilation phase, enabling modularity and flexible code configuration for different build environments.

NODIA

198. (B) Class C can access the private and protected members of class P.
In C++, the friend keyword is used to grant a non-member function or another class full access to the private and protected members of a class. When class C is declared as a friend of class P, it means that members of C can directly interact with P's internal, hidden data. This mechanism should be used sparingly, as it bypasses standard encapsulation principles.

199. (A) (a), (b), (c), and (e) only
Web development relies on several core technologies working together. HTML provides the structural content, while CSS manages the visual presentation. JavaScript enables interactive client-side behavior, and DHTML combines these elements for dynamic updates without page reloads. Conversely, XML is a data

for language designers, ensuring that high-level code can be mapped efficiently regardless of the specific underlying hardware. This abstraction layer bridges the gap between complex source languages and varied machine architectures, facilitating predictable and uniform software execution environments.

252. (A) Lexical analysis, syntax analysis, semantic analysis, and code generation.
The language translation process involves several distinct stages to convert high-level source code into machine-readable instructions. Key phases include lexical analysis for token recognition and syntax analysis for structure validation. Furthermore, semantic analysis ensures logical consistency, while code generation produces the final executable. These sequential steps provide a rigorous and reliable framework for interpreting and translating various programming languages.

253. (C) Binding time
Binding time represents the specific moment when various attributes, such as data types or memory addresses, are associated with program entities like variables. When this association occurs during the compilation phase, it allows the compiler to perform type checking and optimization before execution. This concept is fundamental to understanding how different programming languages manage their resources and ensure type safety.

254. (D) Concepts are the fundamental ideas that are organized into broader programming paradigms.
Programming language concepts encompass the fundamental ideas that define how a language operates, including its data typing and control structures. These foundational concepts are categorized into broader paradigms, which represent distinct schools of thought regarding software development. Therefore, paradigms serve as organizational frameworks that group specific conceptual approaches, such as procedural or object-oriented methods, into cohesive and structured programming models.

255. (C) Lexical analysis
Formal transition models provide a theoretical framework for automating the recognition of patterns within source code. Finite automata are specifically employed during the lexical analysis phase to identify and categorize individual tokens like keywords and identifiers. This application ensures that the initial translation stage is governed by rigorous mathematical principles, leading to predictable and accurate identification of basic language units.

256. (B) It represents a single, indivisible value.
Scalar data types are defined as basic building blocks that represent a single, atomic value. Unlike complex structures, these types are indivisible from the perspective of the programmer and hold fundamental information such as integers, characters, or Booleans. Their primary characteristic is the representation of one discrete value at a time, forming the basis for constructing more complex and nested data structures.

257. (D) A type constructed from one or more elementary data types to form a collection.
Composite data types are structural constructs created by grouping one or more elementary or scalar data types into a single unit. These types are designed to represent collections of related information, which can be either homogeneous or heterogeneous in nature. By organizing simpler data components into broader structures, composite types facilitate sophisticated data modeling and efficient management of complex information sets.

258. (A) An object is a runtime instance of a data type.
The relationship between a data type and an object is defined by their roles during program definition and execution. A data type serves as a blueprint that specifies storage requirements and permissible operations. An object, conversely, is a specific runtime instance of that data type that occupies actual memory and holds a distinct value. This distinction is critical for effective memory management.

259. (C) Composite, because it is a collection of elements.
An array is categorized as a composite data type because it functions as a collection of multiple elements rather than a single atomic value. In the case of an array of integers, it groups together several scalar integer values into a homogeneous structure. This grouping allows programmers to

Downloaded from www.pdf.tube

member in a single, concise step. It is a more direct alternative to using the combination of the dereference operator and the dot operator, which would be written as (*p).data.

209. (C) A is true but R is false.
XML is a versatile markup language designed to describe and store data independently of how it will be displayed. Unlike HTML, which includes tags for visual presentation, XML tags only describe the meaning and structure of the content. To display XML data, separate instructions like CSS or XSLT are required. Therefore, XML is not a self-contained presentational format but a pure data representation tool.

210. (A) A language that is typically interpreted within another program or environment to automate tasks.
Scripting languages are specialized programming languages that are generally interpreted rather than compiled into standalone machine code. They are often used within a larger application or host environment to automate repetitive tasks, manage system configurations, or add interactivity to web pages. Common examples include JavaScript, Python, and Bash, all of which provide a high-level, flexible way to control and extend software systems.

211. (C) A function to define the position and orientation of the camera or viewer.
A viewing function in a graphics API like OpenGL is used to establish the observer's perspective within a 3D scene. This typically involves specifying the camera's location, the point it is looking at, and its orientation in space. By defining these parameters, the function sets the viewing transformation, which is essential for correctly projecting and rendering the scene from the viewer's specific viewpoint.

212. (D) An ellipsoid
An ellipsoid is a primary example of a quadric surface, which is a three-dimensional shape defined by a general second-degree algebraic equation. These surfaces are smooth and continuous, making them ideal for modeling various curved objects in computer graphics with high mathematical precision. Other examples of quadric surfaces include spheres, cylinders, and cones, all of which are defined by similar second-order polynomial equations.

213. (A) Both A and R are true and R is the correct explanation of A.
In the C language, strings are not a separate data type but are instead represented as arrays of characters. The end of each string is marked by a special null character, written as \0. This convention is critical because standard library functions for string manipulation, such as copying or measuring length, rely on finding this null terminator to identify where the character data ends.

214. (A) P-2, Q-4, R-1, S-3
C++ provides keywords for managing class members and memory. The static keyword creates members shared by all objects, while a const member function is guaranteed not to modify data. For dynamic memory, new is used to allocate space and call the constructor, while delete deallocates that memory and triggers the destructor. These tools are fundamental for building efficient and reliable object-oriented software systems.

NODIA

215. (B) (a), (b) and (c) only
Preprocessor directives in C are instructions handled before the actual compilation begins. Common examples include #include for inserting file contents, #define for creating macros, and #ifdef for controlling conditional compilation. These directives are always prefixed with a hash symbol. While virtual is a keyword in C++, it is used as a function modifier and is not a preprocessor directive, as it is handled during compilation.

216. (A) Polymorphism means one name, one form.
Polymorphism literally means many forms and refers to the ability of a single interface to represent different underlying forms or behaviors. This allows objects of different classes to respond to the same message in unique ways. Other core principles include encapsulation for data hiding, inheritance for code reuse, and abstract classes which cannot be instantiated directly, all of which are accurately described in the object-oriented model.

217. (C) cerr

time without needing to send requests back to the server, resulting in a more fluid user experience.

235. (B) Ideal for line drawing but unsuitable for realistic shading.
Random-scan displays are specialized systems that direct the electron beam to trace individual lines and characters directly on the screen. While this approach results in exceptionally smooth and precise line drawings, it is not capable of filling areas with the continuous tones required for realistic shading and complex photographic images. Therefore, these displays were primarily used in early computer-aided design and scientific visualization applications.

236. (B) Multiplying the individual transformation matrices.
A composite transformation matrix is created by the mathematical multiplication of multiple individual transformation matrices, such as those for translation, rotation, and scaling. By performing this multiplication beforehand, a single matrix is produced that represents the combined effect of all the operations. This allows the graphics system to apply the entire sequence of transformations to each point in a single step, significantly improving performance.

237. (C) The curve must pass through the first and last control points.
B-spline curves provide several powerful properties for geometric modeling, including local control and a degree that is independent of the total number of control points. While the curve generally follows the shape of its control polygon, a standard B-spline does not necessarily pass through its first and last control points. This distinguishes it from Bezier curves, which always interpolate their starting and ending points.

238. (A) Both Statement I and Statement II are true.
C++ allows for the declaration of constant pointers in different ways. Using const int *p creates a pointer to a constant integer, meaning the value being pointed to cannot be changed through that pointer. Conversely, int * const p declares a pointer that is itself constant, meaning the address it stores cannot be modified once initialized. Understanding these distinctions is critical for writing safe and correct code.

239. (A) The smallest unit of a program, such as a keyword, operator, or identifier.
In C, a token is the most basic lexical building block that the compiler can recognize. During the initial stage of translation, the source code is partitioned into tokens, which include keywords, identifiers, operators, constants, and string literals. These tokens form the foundation for the program's structure and logic, serving as the input for the subsequent syntax analysis and code generation phases.

240. (A) (a), (b), and (d) only
C++ uses several mechanisms to manage objects and their data. The this pointer uniquely identifies the current instance within member functions, and static members are shared across all instances of a class. To ensure proper cleanup in an inheritance hierarchy, destructors should be declared virtual in the base class. However, constructors cannot be virtual, as they are responsible for creating the object's specific type.

241. (B) To determine which pixels on a raster display should be illuminated to represent a line segment.
Line drawing algorithms are fundamental for rasterizing geometric shapes. Their primary role is to identify the set of individual pixels on a discrete display grid that best approximate a continuous line segment between two specified endpoints. Algorithms like DDA and Bresenham's use incremental calculations to efficiently select these pixels, ensuring that lines appear smooth and accurate while minimizing the total computational overhead required.

242. (A) It shifts points vertically, proportional to their x-coordinate.
In a 2D shear transformation along the y-axis, the shy parameter determines how much each point's vertical coordinate is shifted. This shift is directly proportional to the point's horizontal distance from the origin. This operation distorts the object by slanting it vertically, creating a sliding effect that changes the shape while preserving its area and original width, which is fundamental for creating slanted graphical effects.

243. (C) (c)
Spline surfaces, particularly B-spline and NURBS surfaces, are the industry standard for modeling smooth, complex 3D shapes. Unlike simple

Click Here to Buy book on Amazon

226. (A) It transforms a 3D view volume into a canonical view volume, typically a cube.
In the 3D viewing pipeline, a general projection transform is used to map the visible three-dimensional space, known as the view frustum, into a standardized or canonical view volume. This canonical volume is typically a unit cube. This transformation simplifies the subsequent stages of clipping and rendering by creating a uniform and mathematically convenient space for all projected objects, regardless of the original viewing parameters.

227. (C) Gouraud shading interpolates color intensities, while Phong shading interpolates surface normal vectors.
The primary distinction between these two shading methods lies in what values are interpolated across the polygon. Gouraud shading calculates lighting at the vertices and interpolates the resulting colors, which can miss small specular highlights. Phong shading interpolates the surface normals across the polygon and then calculates the lighting at every individual pixel, providing a much more accurate and realistic representation of highlights and reflections.

228. (B) (b)
When an object of a derived class is created, the base class constructor is always executed first, followed by the derived class constructor. This order ensures that the foundational parts of the object are properly initialized before the more specialized parts are added. This sequence is maintained through all levels of an inheritance hierarchy, providing a consistent and logical initialization process for complex object structures.

229. (C) Debugging
The translation process in a compiler consists of several formal stages, including lexical analysis to identify tokens, syntax analysis to verify structure, and code generation to produce the final executable. Debugging, however, is not a stage of the translation itself but a separate developer activity aimed at finding and fixing errors in the source code or the program's behavior during or after development.

230. (A) P-3, Q-4, R-2, S-1
Fundamental programming concepts are essential for understanding how languages work. Type checking ensures correct data usage, while binding time refers to when associations between entities and attributes are established. A paradigm is the fundamental model or style of programming used. Syntax provides the formal grammatical rules that dictate how symbols must be combined to create a valid program that can be parsed and executed.

231. (D) -> (arrow)
In the C language, the arrow operator is the specialized syntax used for accessing members of a structure through a pointer. It serves as a convenient shorthand that both dereferences the pointer and accesses the specified member in one operation. This is equivalent to using the combination of an asterisk and a dot, but provides a much cleaner and more readable way to handle pointers.

NODIA

232. (A) Encapsulation
Encapsulation is the core object-oriented concept that involves bundling an object's state and behavior within a single unit, usually a class. By using access specifiers to make data members private, the internal state is hidden from outside code. Interaction with the object is then restricted through its public methods, which ensures data integrity and prevents unauthorized or unintended modification of the object's internals.

233. (C) It allocates memory from the heap for an object and calls its constructor.
The new operator in C++ is used for dynamic memory management. It performs two critical tasks: first, it allocates the required amount of memory from the heap or free store for an object; second, it automatically triggers the appropriate constructor to initialize that object. This ensures that the dynamically created object is properly set up and ready for use within the program's logic.

234. (A) Scripting (e.g., JavaScript)
In web programming, client-side scripting is the technology used to make web pages interactive and responsive to user actions. JavaScript is the most common example, running directly within the user's browser. Unlike server-side components like Servlets, client-side scripts can modify the page's content and appearance in real-

representation language, not primarily designed for graphical display, making it distinct from these other presentational and behavioral web technologies.

200. (A) A parametric surface that is an extension of the Bezier curve, defined by a mesh of control points.
A Bezier surface represents a three-dimensional parametric surface derived from the principles of Bezier curves. It is defined using a bidirectional mesh or grid of control points that influence its shape without necessarily passing through all of them. This mathematical model provides a smooth, continuous representation of complex surfaces, making it a standard tool in computer-aided design and modeling.

201. (C) $\theta 1 + \theta 2$
In two-dimensional geometric transformations, the composition of two successive rotations around the same origin is equivalent to a single rotation by the sum of their individual angles. Mathematically, this property stems from the trigonometric identities used in the rotation matrices. This additive nature makes it easy to combine multiple rotational movements into a single transformation, simplifying the overall calculations in a graphics system.

202. (A) To determine which surfaces are visible to the viewer at each pixel location.
The Z-buffer algorithm, also known as the depth-buffer algorithm, is a common technique for visible-surface detection in 3D graphics. It stores the depth, or z-coordinate, of the closest object at each pixel location. When a new polygon is rendered, its depth is compared with the value in the buffer; if the new point is closer, the buffer and the color at that pixel are updated.

203. (B) (b)
The Singleton design pattern is used to ensure that a class has only one instance and provides a global point of access to it. This is typically achieved by making the constructor private and providing a static method that returns the single instance of the class. This pattern is useful for managing shared resources, such as database connections, loggers, or configuration settings, throughout a program.

204. (A) (i) and (iii) are equally parallel, followed by (ii)
Projections are classified by how they represent three-dimensional space. Orthographic and axonometric projections are both types of parallel projections, where lines that are parallel in space remain perfectly parallel in the image. Perspective projection, however, is a non-parallel projection where parallel lines converge toward vanishing points. Therefore, orthographic and axonometric projections are more parallel than the perspective projection model, which simulates realistic depth.

205. (B) P-3, Q-1, R-2, S-4
Core object-oriented concepts provide the framework for sophisticated software design. Abstract classes act as non-instantiable templates for inheritance, while virtual functions enable dynamic behavior by being resolved at runtime. Templates allow for the creation of generic, type-independent code. Constructors are special methods dedicated to the creation and initialization of objects, ensuring they start their lifecycle in a valid and predictable state.

206. (A) Both Statement I and Statement II are true.
Binding in C++ can be static or dynamic. Late binding is usually implemented using virtual function tables, which allow the program to select the correct function version at execution time based on the object's actual type. Static binding is resolved during compilation, which is more efficient because it avoids the overhead of a runtime lookup. Both mechanisms are essential for balancing performance and flexibility.

207. (A) fopen()
The fopen() function is a standard library routine in C used to establish a connection between a program and a physical file on disk. It takes the file name and an access mode, such as read or write, as arguments and returns a file pointer. This pointer is then used by other input/output functions to manage the actual reading or writing of data during program execution.

208. (B) p->data
In C++, the arrow operator is the standard way to access a member of an object when you have a pointer to that object. This operator combines dereferencing the pointer and accessing the

Downloaded from www.pdf.tube

polygon meshes or quadric surfaces, splines offer precise mathematical control and the flexibility of local modification. This allows designers to create highly detailed and aesthetically pleasing surfaces with continuous curvature, which is essential in automotive design, aerospace engineering, and advanced character animation for digital media.

244. (D) polymorphism : The ability to take more than one form
Polymorphism is a core tenet of object-oriented programming that allows a single name or interface to represent multiple different forms or behaviors. This provides flexibility by enabling the same code to interact with objects of different types, with the specific implementation being determined at runtime. Other terms like fopen, argv, and struct have specific, different roles in C programming unrelated to this fundamental object-oriented concept.

245. (C) A is true but R is false.
Templates are a powerful mechanism in C++ that enables generic programming by allowing code to be written with placeholder types. However, templates are expanded by the compiler during the compilation phase, not at runtime. This means that type checking is performed statically, ensuring high performance and type safety. The compiler generates specialized versions of the template code for each specific data type used in the program.

246. (B) To efficiently accept or reject a line segment against a rectangular clipping window.
The Cohen-Sutherland algorithm is a classic technique used in computer graphics to determine whether a line segment lies within, outside, or partially crosses a rectangular clipping window. By assigning bitwise outcodes to the line's endpoints, the algorithm can quickly identify segments that are completely visible or invisible. For lines that are partially visible, it calculates the intersection points with the window boundaries to clip the segment.

247. (D) Perspective projection makes objects appear smaller as their distance from the viewer increases, while parallel projection does not.
The fundamental difference between these two projections lies in their treatment of distance. Perspective projection mimics human vision by making objects that are further away appear smaller, with parallel lines converging at vanishing points. In contrast, parallel (orthographic) projection maintains the original size and proportions of objects regardless of their distance from the viewer, which is why parallel lines always remain perfectly parallel in the projected image.

248. (A) P-3, Q-4, R-1, S-2
C functions are organized into standard header files based on their primary roles. The stdio.h file contains input/output functions, while stdlib.h provides memory management and utility routines. String manipulation functions are located in string.h, and mathematical operations are declared in math.h. By including the correct header, programmers ensure that the compiler has the necessary declarations to use these standard library tools correctly within their source code.

NODIA

249. (C) A class that has at least one pure virtual function, and cannot be instantiated.
An abstract class in C++ serves as a foundational template for other classes and cannot be used to create objects directly. It must contain at least one pure virtual function, which is a member function declared with the = 0 syntax. This design pattern ensures that derived classes must provide their own specific implementations for these functions, promoting a consistent interface while preventing the use of incomplete classes.

250. (D) Statements (a), (b), (c), (d), and (e) are correct.
The new and delete operators in C++ provide a robust framework for dynamic memory management. While new handles memory allocation on the heap and constructor invocation, delete manages deallocation and destructor calls. These operators can also be overloaded to implement custom memory strategies, offering programmers granular control over resource lifecycle and efficiency within complex software applications and object-oriented systems.

251. (B) To offer an abstract, consistent target machine, simplifying the translation process.
A virtual computer functions as an abstract machine designed to simplify the language translation process. It provides a consistent target

In C++, the cerr stream is the standard choice for outputting error messages. It is unbuffered by default, which ensures that any information sent to it is displayed immediately on the screen or in a log file. This is particularly important for debugging and handling critical system failures, where you want to ensure the error data is captured and presented even if the program terminates abruptly.

218. (A) [[1, 0, tx], [0, 1, ty], [0, 0, 1]]
The homogeneous matrix representation for a 2D translation uses a 3x3 matrix where the translation factors are placed in the last column. When a point in homogeneous coordinates is multiplied by this matrix, the x and y coordinates are shifted by tx and ty respectively. This representation is essential because it allows translation to be treated as a multiplicative operation, consistent with other geometric transformations.

219. (A) (a) and (b) only
Basic file handling and error management in C++ are primarily addressed through the use of file stream classes and exception handling. The ofstream and ifstream classes provide the core functionality for reading from and writing to files. Additionally, try and catch blocks offer a structured way to handle potential runtime errors that may occur during these input/output operations, ensuring program robustness and reliability.

220. (B) union
A union is the most memory-efficient C construct for storing mutually exclusive data of different types. Because all members of a union share the same memory location, only enough space for the largest member is allocated. This is ideal when only one piece of information needs to be active at a time, allowing for compact data representation while still supporting a variety of different data formats.

221. (C) A derived class can access the private members of its base class.
In C++ inheritance, private members of a base class are never directly accessible to any other class, including its own derived classes. Access is restricted to public and protected members. While a derived class inherits all properties, it can only interact with private data through inherited public or protected functions. This strict rule ensures strong encapsulation and protects the internal integrity of the base class.

222. (D) Built-in graphical rendering capabilities.
XML is a powerful language for representing and exchanging data in a platform-independent and structured way. Its primary strengths are its extensibility and the separation of data from its presentation. However, XML itself does not have built-in graphical rendering capabilities. Displaying XML content requires additional technologies like CSS or XSLT to define how the structured data should be visually presented in a web browser or other software.

223. (A) (a), (b), and (c) only
Fundamental geometric transformations in computer graphics consist of translation, rotation, and scaling. These operations are the basic building blocks used to manipulate the position, orientation, and size of objects within a scene. Clipping, while essential for rendering, is not a geometric transformation but a process used to identify and discard parts of an image that fall outside the defined viewing window or volume.

224. (B) (a), (c), d, and (e) only
Computer graphics relies on several core principles. Raster-scan systems use frame buffers to store pixel data, while the mid-point circle algorithm provides efficient integer-based calculations. Clipping is necessary to restrict the rendered image to the viewport, and the viewing pipeline manages the overall transformation from world to device coordinates. However, random-scan systems do not scan line-by-line; they draw geometric primitives directly.

225. (B) A region on the display device where the window is mapped.
A viewport is a designated rectangular area on a computer screen or display device where the contents of a viewing window are rendered. While the window defines what part of the world-coordinate scene is selected for display, the viewport determines the specific location and size of that image on the physical output device. This mapping process involves scaling and positioning the selected scene correctly.

Click Here to Buy book on Amazon

manage large sets of related data under a single identifier using indexed access methods.

260. (A) The properties of its constituent types and its structural rules.
The properties of a composite data type, like a C structure, are determined by the characteristics of its individual components and the specific rules governing its internal arrangement. These rules define how the constituent scalar or composite types are organized and accessed within memory. Consequently, the overall behavior of the structure is a direct result of its members and its structural definition.

261. (D) To store a memory address, enabling direct memory manipulation.
In the C language, pointers are specialized variables designed to store the memory addresses of other data entities. By holding these addresses, pointers allow for direct manipulation of memory and efficient management of complex structures like linked lists or dynamic arrays. This capability provides low-level control over hardware resources, which is a defining feature for system-level programming and resource optimization.

262. (B) Structures store heterogeneous data, while arrays store homogeneous data.
Structures and arrays are both composite types in C but differ in the nature of the data they store. An array is a collection of homogeneous elements, meaning all items must share the same data type. In contrast, a structure is a heterogeneous collection, allowing the grouping of diverse data types into a single unit to represent complex real-world entities effectively.

263. (C) It allows different data types to be stored in the same memory location.
A union is a unique data type in C that enables different variables of various types to share the exact same memory location. While it can store multiple types, only one member can contain a value at any given time. This feature is particularly useful for memory conservation and for interpreting the same set of bits in multiple different ways during execution.

264. (D) Command line arguments
Command line arguments provide a standard mechanism in C for passing parameters to a program exactly at the moment it begins execution. This feature allows users to provide input or configuration settings directly from the terminal. By accepting these external parameters, programs become more flexible and can adjust their behavior without needing to be recompiled for every different use case.

265. (A) Identifiers, keywords, and operators.
Tokens represent the smallest meaningful units of a C program that the compiler recognizes during the initial phase of translation. These fundamental lexical components include identifiers for naming variables, keywords for language commands, and operators for performing calculations. Every C program is essentially a sequence of these tokens, which are analyzed and structured into higher-level code by the compiler.

266. (B) Instantiation
Instantiation is the specific process in object-oriented programming where a concrete object is created based on the blueprint provided by a class. During this process, memory is allocated for the object's data members as defined in the class structure. This transformation from a conceptual class definition to a functional runtime instance allows the program to interact with data.

267. (B) Encapsulation
Encapsulation is a core principle of object-oriented programming that involves bundling data and the methods that operate on it into a single unit. This mechanism restricts direct access to the internal state of an object, requiring all interactions to occur through a defined public interface. This approach protects data integrity, promotes modularity, and simplifies the maintenance of complex software systems.

268. (A) To serve as a base class that cannot be instantiated but defines an interface for subclasses.
An abstract class serves as a foundational base class that is specifically designed not to be instantiated on its own. Its primary purpose is to define a common interface and shared behaviors for a group of related subclasses. By including abstract methods, it mandates that any derived class must provide its own implementation,

Downloaded from www.pdf.tube

ensuring structural consistency across a class hierarchy.

269. (C) Code reuse by allowing a new class to acquire properties from an existing class.
Inheritance is a powerful mechanism in object-oriented programming that primarily facilitates code reuse by allowing new classes to acquire properties and behaviors from existing ones. By establishing a hierarchical relationship, a derived class can extend or modify the functionality of its base class. This minimizes redundancy and enables the development of organized and scalable software architectures through specialized class structures.

270. (D) An abstract class
An abstract class is the most appropriate feature for defining a common interface among a group of related classes. It allows a programmer to specify which methods must be implemented by every subclass without providing the code for them. This ensures that all objects within that hierarchy adhere to a consistent contract, which is essential for maintaining polymorphism and structure.

271. (C) To implement runtime polymorphism within a class hierarchy.
Virtual functions are a critical component in C++ for implementing runtime polymorphism within a class hierarchy. They allow a program to decide which implementation to execute at runtime based on the actual type of the object being pointed to. This dynamic binding enables flexible design, where a single interface supports multiple behaviors across different derived classes.

272. (B) Constructors and destructors
Constructors and destructors are specialized member functions in C++ that manage the lifecycle of objects. A constructor is automatically called to initialize an object's data members upon its creation, ensuring it starts in a valid state. Conversely, a destructor is executed when an object is destroyed, providing a mechanism to release resources and perform necessary cleanup to prevent memory leaks.

273. (D) Templates
Templates in C++ provide a robust mechanism for generic programming, allowing the creation of functions and classes that can operate with any data type. By using placeholders for types, a single template definition can be reused for integers, floats, or custom objects. This feature significantly reduces code duplication and enhances the flexibility and maintainability of large-scale software development projects.

274. (A) By supporting the organization of code into multifile programs.
C++ facilitates the development of large and modular applications by allowing programmers to organize their source code into multiple files. This multifile approach enables different modules to be developed and compiled independently, which simplifies management and improves build efficiency. By separating concerns into distinct files, teams can collaborate more effectively on complex software systems while maintaining clear code structure.

275. (C) Overloading
Overloading is a feature in C++ that allows multiple functions or operators to share the same name but have different implementations. The compiler distinguishes between these versions based on the number or type of arguments provided during a call. This capability enhances code readability and allows similar operations to be performed on different data types using a consistent and intuitive naming convention.

276. (A) HTML is focused on the presentation of content, while XML is focused on describing data.
HTML and XML serve fundamentally different purposes within web development. HTML is primarily a markup language designed for the presentation of content, defining how text and images appear on a web page. In contrast, XML focuses on describing and structuring the data itself. While HTML has a fixed set of tags, XML allows users to define custom tags.

277. (D) Servlet
A Servlet is a specialized Java component that operates on a web server to handle incoming requests and generate dynamic responses. It serves as a middle layer between the client's browser and server-side resources like databases.

NODIA

By executing logic on the server, Servlets enable the creation of interactive web applications that can provide customized content based on user input or state.

278. (B) A combination of technologies used to dynamically modify web page content.
Dynamic HTML, or DHTML, is a collection of technologies, including HTML, CSS, and scripting languages, used to create interactive and animated web pages. It is not a standalone language but rather a method for modifying the content and appearance of a page after it has been loaded into a browser. This allows for real-time updates without requiring a full page refresh.

279. (C) Applets are executed within a web browser on the client-side, while Servlets run on the web server.
The primary difference between an Applet and a Servlet lies in where they execute. An Applet is a small Java application that is downloaded to the client's machine and runs within a web browser. In contrast, a Servlet is a server-side component that remains on the web server to process requests and generate responses, providing centralized logic and data management.

280. (D) Scripting
Client-side interactivity is most directly achieved through scripting languages like JavaScript that run inside the user's web browser. These scripts can manipulate the page content in real-time, perform calculations, and validate user input before it is sent to a server. This immediate feedback enhances the user experience by making web pages feel responsive and interactive without constant server communication.

281. (A) Raster-scan sweeps the entire screen row-by-row, while random-scan directs the beam only where the image is.
Raster-Scan and Random-Scan systems differ in how they direct an electron beam to create images. A Raster-Scan system sweeps the beam across the entire screen row by row, illuminating pixels to form a grid-based image. Conversely, a Random-Scan system directs the beam only to specific areas where lines and shapes are to be drawn, behaving more like a pencil.

282. (B) Determine which pixels to illuminate to approximate a line on a raster display.
Line drawing algorithms, such as the Bresenham's algorithm, are essential for representing geometric shapes on a pixel-based display. Their primary role is to calculate which specific pixels in a raster grid should be turned on to best approximate a continuous straight line. This ensures that the digital representation of a line appears as smooth and accurate as possible on the screen.

283. (B) Boundary-Fill algorithm
The Boundary-Fill algorithm is designed to fill an irregularly shaped area that is defined by a specific border color. It works by starting from a specified interior point and recursively coloring adjacent pixels until the boundary color is reached. This makes it ideal for coloring regions in graphics applications where the shape might be complex but has a clearly defined edge.

284. (A) By processing the polygon's area on a horizontal, row-by-row basis.
The Scan Line Polygon Fill algorithm operates by processing a polygon on a horizontal, row-by-row basis. It calculates the intersections of each scan line with the edges of the polygon to determine which segments of the line fall inside the shape. These interior segments are then filled with color, allowing for the efficient rendering of complex polygonal areas on displays.

285. (C) It represents images as a grid of pixels and uses a Raster-Scan approach.
Modern graphics monitors are characterized by their use of the Raster-Scan approach, where images are represented as a dense grid of pixels. In this system, the display is refreshed by scanning an electron beam or updating electronic signals across the screen in a consistent pattern. This method allows for the display of complex, high-resolution images and is the industry standard.

286. (D) They allow all geometric transformations, including translation, to be represented as matrix multiplications.
Homogeneous coordinates are introduced in 2-D graphics to create a unified mathematical framework for all geometric transformations. By representing points as three-element vectors, complex operations like translation, which is

Click Here to Buy book on Amazon

CHAPTER 04

Database Management Systems

1. Which of the following best defines a 'database schema'?
(A) The collection of data stored in the database at a particular moment.
(B) The overall design or description of the database structure.
(C) The language used to query and manipulate the database.
(D) The physical storage mechanism for the database files.

2. Consider the following statements and select the correct option.
Statement I : In a centralized DBMS architecture, the DBMS software and data reside on a single computer.
Statement II : A client/server architecture divides the system into two main components: a server that provides database functionality and clients that request services.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

3. Which of the following is an example of a 'data model'?
(A) Structured Query Language (SQL)
(B) Oracle Database Management System
(C) The Entity-Relationship Model
(D) A database trigger mechanism

4. Match the items in List I with the descriptions in List II and select the correct option.
List I: (Schema Level)
P. Internal Schema
Q. Conceptual Schema
R. External Schema
S. Data Independence

List II: (Description)
1. Describes the part of the database that a particular user group is interested in.
2. The ability to modify a schema definition in one level without affecting a schema definition in the next higher level.
3. Describes the physical storage structure of the database.
4. Describes the structure of the whole database for a community of users.

(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-2, S-1
(C) P-3, Q-1, R-4, S-2
(D) P-1, Q-2, R-3, S-4

5. What is the term for the data stored in a database at a particular moment in time?
(A) Schema
(B) Model
(C) Instance
(D) Interface

6. Which of the following is NOT a type of database language?
(A) Data Definition Language (DDL)
(B) Data Processing Language (DPL)
(C) Data Manipulation Language (DML)
(D) Data Control Language (DCL)

7. The ability to change the internal schema without having to change the conceptual schema is known as:
(A) Logical Data Independence
(B) Architectural Data Independence
(C) Conceptual Data Independence
(D) Physical Data Independence

8. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : The three-schema architecture provides data independence.
Reason (R) : It separates the user view of the

NODIA

the smooth, complex skins of 3-D objects. These representations are vital in engineering and design for developing continuous, high-quality surfaces that are necessary for manufacturing and realistic rendering.

296. (D) To simulate the interaction of light with surfaces to determine their color and brightness.
An illumination model is used in 3-D graphics to simulate the physical interaction of light with various surfaces. It calculates the intensity and color of light at specific points based on the properties of light sources and the material of the object. This simulation is essential for creating realism, as it produces the highlights, shadows, and shading necessary to represent depth.

297. (A) Perspective projection
Perspective projection is a transformation technique used in the 3-D viewing pipeline to create a sense of depth realism on a 2-D screen. It operates on the principle that objects further away from the viewer should appear smaller than those that are closer. This mimicry of human vision is fundamental for rendering realistic environments where spatial relationships and distance are clearly perceived.

298. (B) Compute colors across a polygon's surface for a smooth appearance.
Polygon rendering methods, such as Gouraud and Phong shading, are designed to compute color values across the surface of a polygon. Instead of coloring a facet uniformly, these methods interpolate values to create a smooth transition between vertices. This technique masks the flat edges of a polygonal mesh, making a faceted object appear as a smooth, continuous surface in the final image.

299. (B) To remove parts of objects that lie outside the defined view volume.
In the 3-D viewing pipeline, clipping is the process of identifying and removing parts of objects that fall outside the three-dimensional view volume. This step ensures that the graphics system only processes geometry that is actually visible from the current camera position. By discarding hidden or external data, clipping optimizes the rendering process and prevents visual artifacts at the screen edges.

300. (A) Viewing pipeline
The viewing pipeline is the complete sequence of coordinate transformations and operations that convert 3-D object representations into a final 2-D image on a display. This complex process includes mapping world coordinates to view coordinates, applying projections, performing clipping, and rendering surface colors. It provides a structured framework for ensuring that a 3-D scene is accurately depicted from any given viewpoint.

www.nodia.in

Downloaded from www.pdf.tube

34. A programmer wants to prevent SQL injection by ensuring that user input is treated as data and not as part of the SQL command.
What is the most appropriate technique?
(A) Using prepared statements with parameterized queries
(B) Using dynamic SQL concatenation
(C) Blacklisting specific characters like ‘ and --
(D) Limiting the length of user input fields

35. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : Normalization is a crucial step in relational database design.
Reason (R) : It helps to minimize data redundancy and reduce the likelihood of update anomalies.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

36. Which of the following is not a DML (Data Manipulation Language) command in SQL?
(A) `INSERT`
(B) `UPDATE`
(C) `DELETE`
(D) `CREATE`

37. What is the purpose of the `UNIQUE` constraint in SQL?
(A) It uniquely identifies each record in a database table and must contain unique values.
(B) It prevents any NULL values from being inserted into a column.
(C) It ensures that all values in a column are different, but allows one NULL value.
(D) It is used to create a relationship between two tables.

38. Which type of database object is automatically invoked by the DBMS when a specified data manipulation event occurs?
(A) Stored Procedure
(B) Trigger
(C) View
(D) Function

39. Consider a database for an online store.
When a new order is placed in the `Orders` table, the inventory count for the corresponding product in the `Products` table must be decreased.
What is the most suitable database mechanism to automate this action?
(A) A view on the `Products` table
(B) A stored procedure called manually by the application
(C) An SQL injection check
(D) A trigger on the `Orders` table

40. Consider the following statements regarding SQL (Structured Query Language).
a. SQL is a declarative language.
b. `CREATE`, `ALTER`, and `DROP` are Data Definition Language (DDL) commands.
c. `GRANT` and `REVOKE` are Data Control Language (DCL) commands.
d. `SELECT`, `INSERT`, and `UPDATE` are Data Manipulation Language (DML) commands.
e. SQL is case-sensitive for keywords like SELECT and FROM.
Which of the above statements are correct?
(A) a, b, and e only
(B) c, d, and e only
(C) a, b, c, and d only
(D) a, b, and d only

41. A functional dependency $X \rightarrow Y$ holds in a relation R if:
(A) For any two tuples t1 and t2 in R that have $t1[Y] = t2[Y]$, they must also have $t1[X] = t2[X]$.
(B) The attribute set X is a candidate key of the relation R.
(C) The attribute set Y is functionally dependent on all attributes in R.
(D) For any two tuples t1 and t2 in R that have $t1[X] = t2[X]$, they must also have $t1[Y] = t2[Y]$.

42. A relation is in First Normal Form (1NF) if and only if:
(A) It contains no transitive dependencies.
(B) The domain of each attribute contains only atomic values.
(C) It contains no partial dependencies.
(D) Every non-key attribute is fully functionally

database from the physical representation of the database.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) Assertion A is true but the provided Reason R is false.
(D) Assertion A is false but the provided Reason R is true.

9. What is the primary role of a Data Definition Language (DDL)?
(A) To query the database for information retrieval.
(B) To control access and permissions to the database.
(C) To define the database schema and its structure.
(D) To insert, update, and delete data from the database.

10. Which of the following is NOT a type of database language?
(A) Relationships between tables are represented by pointers.
(B) It represents data as a collection of tables.
(C) Each table consists of rows (tuples) and columns (attributes).
(D) Constraints are used to maintain data integrity.

11. Consider the following statements about Codd's rules for a relational database system.
a. The Information Rule states that all information must be stored in tables.
b. The Guaranteed Access Rule ensures every data element is accessible via table name, primary key, and column name.
c. The Systematic Treatment of Null Values rule is a fundamental requirement.
d. The View Updating Rule states that all views that are theoretically updatable must be updatable by the system.
Which of the above statements are correct?
(A) a and b only
(B) a, b, and c only
(C) b, c, and d only
(D) a, b, c, and d

12. In an Entity-Relationship (ER) diagram, what does a double-lined rectangle represent?
(A) A strong entity set
(B) A weak entity set
(C) A relationship set
(D) An attribute

13. Which of the following pairs related to the Relational Model is correctly matched?
(A) Tuple : Column
(B) Attribute : Row
(C) Relation : Table
(D) Domain : Primary Key

14. Identify the chronological sequence of the following database development steps.
a. Conceptual Design
b. Requirements Collection and Analysis
c. Logical Design
d. Physical Design
Which of the following represents the correct sequence?
(A) a, b, c, d
(B) b, a, c, d
(C) b, c, a, d
(D) a, c, d, b

15. Consider the following statements regarding Relational Algebra and Relational Calculus.
a. Relational Algebra is a procedural query language.
b. Relational Calculus is a non-procedural query language.
c. Relational Algebra specifies 'how' to get the result.
d. Relational Calculus specifies 'what' result to obtain without specifying 'how' to obtain it.
e. Both are formally equivalent in expressive power.
Which of the above statements are correct?
(A) a, b, and c only
(B) b, d, and e only
(C) a, c, and d only
(D) a, b, c, d, and e

16. What is the primary purpose of a foreign key constraint in a relational database?
(A) To uniquely identify each record in a table.
(B) To establish and enforce a link between data in two tables.
(C) To ensure that a column cannot have a

Click Here to Buy book on Amazon

(DAC) is based on the identity of the users and/or the groups to which they belong.

Statement II : Mandatory Access Control (MAC) is based on security classifications (or labels) assigned to both users and data objects.

(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

53. An Object-Relational Database (ORDBMS) aims to bridge the gap between which two models?
(A) Hierarchical and Network models
(B) Relational and NoSQL models
(C) Relational and Object-Oriented models
(D) Document and Key-Value models

54. In the context of database transactions, what does a `ROLLBACK` operation do?
(A) It permanently saves all the changes made by the transaction.
(B) It creates a savepoint within the transaction.
(C) It starts a new transaction.
(D) It undoes all the changes made by the transaction.

55. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.

Assertion (A) : The Two-Phase Locking (2PL) protocol ensures serializability.

Reason (R) : In 2PL, a transaction cannot request a new lock after it has released any of its existing locks, which prevents cycles in the precedence graph.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

56. Which of the following is NOT a characteristic feature of an Object-Oriented Database?
(A) Complex objects
(B) Normalization to BCNF
(C) Class hierarchy
(D) Encapsulation

57. Granting and revoking privileges on database objects are functions of which language subset?
(A) Data Control Language (DCL)
(B) Data Definition Language (DDL)
(C) Data Manipulation Language (DML)
(D) Transaction Control Language (TCL)

58. What is a dirty read in the context of concurrency control?
(A) A transaction re-reads a data item and finds that it has been modified by another committed transaction.
(B) A transaction reads a data item that has been deleted by another transaction.
(C) A transaction reads data that has been modified by another transaction that has not yet committed.
(D) A transaction reads data from a table that has no primary key.

59. You are given a relation R(A,B,C,D,E) with primary key AB. You observe the functional dependency C -> D. This violates which normal form?
(A) First Normal Form (1NF)
(B) Second Normal Form (2NF)
(C) Third Normal Form (3NF)
(D) Boyce-Codd Normal Form (BCNF)

NODIA

60. Consider the following database recovery techniques.
a. Log-based recovery
b. Shadow paging
c. Checkpointing
d. Two-phase locking
e. Timestamp ordering

Which of the above are valid recovery techniques?
(A) a, b, and c only
(B) a, b, d, and e only
(C) d and e only
(D) a, b, c, d, and e

61. Which of the following best describes a Temporal Database?
(A) A database designed for very fast, temporary storage of transactional datasets.
(B) A database that only stores the current state

R. Domain Constraint
S. Key Constraint

List II: (Definition)

1. The value of a foreign key must either be a primary key value in another relation or be null.
2. Every value of an attribute must be from its specified set of valid values.
3. No two tuples in a relation can have the same value for the superkey attributes.
4. No primary key value can be null.

(A) P-1, Q-4, R-3, S-2
(B) P-4, Q-1, R-2, S-3
(C) P-4, Q-2, R-1, S-3
(D) P-2, Q-3, R-4, S-1

26. Which SQL statement is used to add a new column to an existing table?
(A) `ALTER TABLE table_name ADD COLUMN new_col_name datatype;`
(B) `ADD COLUMN new_col_name datatype TO table_name;`
(C) `MODIFY TABLE table_name ADD new_col_name datatype;`
(D) `UPDATE TABLE table_name ADD COLUMN new_col_name datatype;`

27. The `VARCHAR` data type in SQL is used to store:
(A) Fixed-length character strings.
(B) Numerical data with high precision.
(C) Date and time values.
(D) Variable-length character strings.

28. What is the primary purpose of an SQL View?
(A) To physically store a subset of data from one or more tables.
(B) To provide a virtual table based on the result-set of an SQL statement.
(C) To create a backup copy of a table.
(D) To enforce referential integrity between tables.

29. Which of the following is a potential consequence of an SQL Injection attack?
(A) Improved database query performance.
(B) Automatic enforcement of data integrity.
(C) Unauthorized access to sensitive data.
(D) Enhanced data compression.

30. Consider the following statements regarding SQL triggers.
a. A trigger is a stored procedure that is automatically executed in response to certain events.
b. Triggers can be defined to execute before or after an INSERT, DELETE, or UPDATE operation.
c. Triggers cannot call stored procedures.
d. Triggers are primarily used to maintain data integrity and enforce complex business rules.

Which of the above statements are correct?
(A) a, b, and d only
(B) a and c only
(C) b, c, and d only
(D) a, b, c, and d

31. Which SQL clause is used to filter results based on an aggregate function?
(A) `HAVING`
(B) `WHERE`
(C) `FROM`
(D) `GROUP BY`

32. Consider the following statements and select the correct option.
Statement I : A stored procedure is a prepared SQL code that you can save, so the code can be reused over and over again.
Statement II : A function in SQL must return a single value or a table, whereas a stored procedure is not required to return a value.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

33. Which of the following SQL commands is used to remove a view from the database?
(A) `DELETE VIEW view_name`
(B) `DROP VIEW view_name`
(C) `REMOVE VIEW view_name`
(D) `ERASE VIEW view_name`

NODIA

normally additive, can be expressed as matrix multiplications. This unification allows multiple transformations to be combined into a single 3x3 matrix, streamlining the computation of composite movements and resizing operations.

287. (C) Multiplying the individual transformation matrices for rotation and scaling.
A composite transform is created by multiplying the individual transformation matrices together in a specific sequence. For example, if an object must be rotated and then scaled, the final transformation matrix is the product of the rotation and scaling matrices. This mathematical property allows the graphics system to apply multiple changes to an object simultaneously using a single combined matrix multiplication.

288. (B) Mapping a selected part of the scene onto an area of the display device.
The window-to-viewport coordinate transformation is a critical step in the viewing pipeline that maps a selected portion of a scene to a specific area on a display device. The window defines what part of the world coordinates is seen, while the viewport defines where that image appears on the screen. This mapping ensures that objects are correctly scaled and positioned.

289. (D) Shear
Shear is a geometric transformation that distorts the shape of an object by shifting its points in a specific direction based on their distance from an axis. This results in a slanting effect, effectively turning a rectangle into a parallelogram. Unlike translation or rotation, shear changes the internal angles of an object, providing a way to represent distortions or perspective effects.

290. (A) To discard portions of objects that are outside the viewing window.
In the context of a 2-D viewing pipeline, a clipping algorithm is responsible for removing any parts of a graphical object that lie outside the designated viewing window. By identifying and discarding these external portions, the algorithm ensures that only the relevant parts of the scene are processed for display. This optimization prevents the rendering of unnecessary data and maintains visual boundaries.

291. (C) Polygon surfaces
Polygon surfaces are the most common scheme for approximating the boundary of a 3-D object using a mesh of flat polygonal facets. By using simple shapes like triangles or quadrilaterals, complex objects can be represented in a way that is computationally efficient for rendering. This method balances the need for detail with the processing limitations of modern graphics hardware and algorithms.

292. (A) To model complex, free-form shapes with smooth curves and surfaces.
Bezier and B-spline representations are primarily used to model smooth, free-form curves and surfaces in computer-aided design and graphics. These spline-based methods provide high continuity and allow for the creation of complex shapes that would be difficult to represent using simple polygons. They are essential for designing objects with aesthetic or aerodynamic requirements, such as car bodies or character models.

293. (D) They offer more localized control over the curve's shape.
A significant advantage of B-spline curves compared to Bezier curves is the localized control they provide over the shape of a curve. In a B-spline, moving a single control point only influences a small segment of the curve rather than its entire length. This characteristic allows designers to make precise adjustments to specific parts of a model without unintentionally altering the whole.

294. (B) Quadric surface
Quadric surfaces are the most efficient and accurate way to represent objects with regular, curved geometries such as spheres, cylinders, or cones. These surfaces are defined by second-degree mathematical equations, which allow for a perfectly smooth representation using minimal data. Compared to polygonal meshes, quadric representations provide a more precise description of these specific shapes while requiring significantly less storage space.

295. (C) Bezier and B-spline surfaces
The extension of spline curve formulations into three dimensions results in the creation of Bezier and B-spline surfaces. These surfaces are defined by a grid of control points and are used to model

Downloaded from www.pdf.tube

dependent on the primary key.

43. Which of the following statements about transaction properties (ACID) is incorrect?
(A) Atomicity: A transaction is an atomic unit of processing; it is either performed in its entirety or not performed at all.
(B) Consistency: A transaction must transform the database from one consistent state to another.
(C) Durability: The changes applied to the database by a committed transaction can be lost in case of a system failure.
(D) Isolation: A transaction should appear as though it is being executed in isolation from other transactions.

44. Which concurrency control technique involves checking for conflicts when a transaction is about to commit?
(A) Two-Phase Locking (2PL)
(B) Timestamp Ordering
(C) Multiversion Concurrency Control
(D) Optimistic Concurrency Control (Validation Based)

45. Match the items in List I with the properties in List II and select the correct option.
List I: (Normal Form)
P. 2NF
Q. 3NF
R. BCNF
S. 1NF
List II: (Property)
1. Eliminates transitive dependencies.
2. For every non-trivial functional dependency $X \rightarrow A$, X must be a superkey.
3. All attributes have atomic domains.
4. Eliminates partial dependencies.
(A) P-1, Q-4, R-3, S-2
(B) P-4, Q-2, R-1, S-3
(C) P-4, Q-1, R-2, S-3
(D) P-3, Q-1, R-2, S-4

46. What is the primary goal of query optimization?
(A) To ensure the query returns a correct result.
(B) To find the query execution plan with the minimum resource usage.
(C) To increase the number of tuples returned by a query.
(D) To parse the SQL syntax correctly.

47. A relation is in Second Normal Form (2NF) if it is in 1NF and:
(A) No non-prime attribute is transitively dependent on the primary key.
(B) Every determinant is a candidate key.
(C) All attributes are functionally dependent only on the primary key.
(D) No non-prime attribute is partially dependent on any candidate key.

NODIA

48. Consider a relation R(A, B, C, D) with functional dependencies $\{A \rightarrow B, B \rightarrow C\}$.
This relation has a:
(A) Transitive dependency
(B) Partial dependency
(C) Multi-valued dependency
(D) Join dependency

49. Which of the following recovery techniques involves restoring the database to a previous consistent state and then re-applying the actions of committed transactions?
(A) Immediate Update
(B) Deferred Update
(C) Checkpointing with Log-Based Recovery
(D) Shadow Paging

50. Consider the following statements about Boyce-Codd Normal Form (BCNF).
a. A relation is in BCNF if for every one of its non-trivial functional dependencies X ? Y, X is a superkey.
b. Every relation in BCNF is also in 3NF.
c. Every relation in 3NF is also in BCNF.
d. BCNF is stricter than 3NF.
Which of the above statements are correct?
(A) a, b, and d only
(B) a, c, and d only
(C) b and c only
(D) a, b, c, and d

51. The Two-Phase Locking protocol for concurrency control consists of which two phases?
(A) Read Phase and Write Phase
(B) Lock Phase and Unlock Phase
(C) Validation Phase and Commit Phase
(D) Growing Phase and Shrinking Phase

52. Consider the following statements and select the correct option.
Statement I : Discretionary Access Control

NULL value.
(D) To validate that the data in a column conforms to a specific format.

17. Which relational algebra operation is used to select tuples from a relation that satisfy a given predicate?
(A) Projection (π)
(B) Selection (σ)
(C) Union ($\cup$)
(D) Cartesian Product ($\times$)

18. Consider the following statements and select the correct option.
Statement I : Tuple Relational Calculus is based on finding tuples for which a predicate is true.
Statement II : Domain Relational Calculus uses domain variables that take on values from an attribute's domain, rather than values for an entire tuple.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

19. Which of the following is NOT one of Codd's 12 rules?
(A) Information Rule
(B) View Updating Rule
(C) Graphical User Interface Rule
(D) Logical Data Independence

20. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : In the relational model, all update operations (insert, delete, modify) on a schema must handle potential constraint violations.
Reason (R) : Relational databases enforce integrity constraints like primary keys, foreign keys, and domain constraints to maintain data accuracy and consistency.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

21. In case of a constraint violation during an INSERT operation, which of the following is the most common default action taken by the DBMS?
(A) Cascade the operation to related tables.
(B) Set the invalid fields to NULL.
(C) Reject the operation.
(D) Log the error and accept the operation.

22. The relational algebra expression π_{A1, A2, ... An} (R) is used for which operation?
(A) Selection
(B) Projection
(C) Join
(D) Division

23. Which statement correctly identifies a difference between Tuple Relational Calculus (TRC) and Domain Relational Calculus (DRC)?
(A) TRC is procedural, whereas Domain Relational Calculus is non-procedural in nature.
(B) TRC is strictly less expressive than Domain Relational Calculus in expressive power.
(C) TRC uses '$\exists$' and '$\forall$' quantifiers, whereas Domain Relational Calculus does not use them.
(D) TRC variables range over tuples, while DRC variables range over domains of attributes.

24. A student is designing an E-R diagram for a university.
A 'Professor' can 'advise' multiple 'Students', but each 'Student' is advised by only one 'Professor'. What is the cardinality of the 'advise' relationship from 'Professor' to 'Student'?
(A) One-to-Many (1:N)
(B) One-to-One (1:1)
(C) Many-to-One (N:1)
(D) Many-to-Many (N:M)

25. Match the items in List I with the definitions in List II and select the correct option.
List I: (Constraint Type)
P. Entity Integrity
Q. Referential Integrity

Click Here to Buy book on Amazon

of data for immediate queries.
(C) A database specifically optimized for managing complex time-zone conversions.
(D) A database that stores data relating to time instances, associating a specific time period with each stored tuple.

62. A database designed to store and query data that is represented in terms of logic-based rules is called a:
(A) Multimedia Database
(B) Deductive Database
(C) Mobile Database
(D) Temporal Database

63. Which technology is central to the structure and exchange of data in Internet Databases?
(A) SQL
(B) Java
(C) XML (eXtensible Markup Language)
(D) HTML (HyperText Markup Language)

64. What is a key challenge in managing Multimedia Databases that is less significant in traditional databases?
(A) Ensuring data atomicity.
(B) Enforcing referential integrity.
(C) Normalizing data to 3NF.
(D) Managing large object sizes and content-based retrieval.

65. Match the Enhanced Data Model in List I with its primary application area in List II and select the correct option.
List I: (Data Model)
P. Geographic Information Systems (GIS)
Q. Genome Data Management
R. Mobile Databases
S. Temporal Databases
List II: (Application Area)
1. Financial analysis and historical record keeping.
2. Managing bioinformatics data like DNA sequences.
3. Handling spatial data and location-based queries.
4. Supporting disconnected operations for handheld devices.
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-2, R-4, S-1
(C) P-3, Q-4, R-2, S-1
(D) P-2, Q-1, R-4, S-3

66. Which characteristic is a defining feature of a Distributed Database?
(A) Data is stored on a single, powerful mainframe computer.
(B) Data is logically interrelated but stored in multiple physically separate locations.
(C) Data is stored in memory for faster access.
(D) Data is unstructured and schemaless.

67. Which of the following is NOT an advantage of a client-server architecture for databases?
(A) Better performance through distributed processing.
(B) Increased security as the server can control access.
(C) Lower hardware costs as clients can be less powerful.
(D) Simpler recovery protocols due to single-site operation.

68. What is the main purpose of Geographic Information Systems (GIS)?
(A) To capture, store, manipulate, analyze, and manage spatial or geographical data.
(B) To manage historical versions of data.
(C) To deduce new facts from a set of rules and facts.
(D) To handle data for mobile applications with intermittent connectivity.

69. In the context of Genome Data Management, what type of data is primarily handled?
(A) Financial transaction records.
(B) Video and audio streams.
(C) Biological sequence and structure data.
(D) Geographical map layers.

NODIA

70. Arrange the following database types in the chronological order of their conceptual development.
a. Relational Database
b. Hierarchical Database
c. Object-Oriented Database
d. NoSQL Database
Which of the following represents the correct sequence?
(A) b, a, c, d
(B) a, b, c, d
(C) b, c, a, d

used for this task?
(A) Data Manipulation Language (DML)
(B) Data Definition Language (DDL)
(C) Data Control Language (DCL)
(D) Transaction Control Language (TCL)

127. Which pair related to the Three-Schema Architecture is incorrectly matched?
(A) Conceptual Level : Describes what data is stored.
(B) Internal Level : Describes how data is stored.
(C) Data Independence : Tightly couples the schema levels.
(D) External Level : Describes a single user's view.

128. What is the main difference between a centralized and a client/server DBMS architecture?
(A) The type of data model used.
(B) The physical distribution of DBMS components.
(C) The query language supported.
(D) The level of data normalization.

129. The operation in relational algebra that returns the columns of a relation is:
(A) π (Projection)
(B) σ (Selection)
(C) ? (Join)
(D) $\div$ (Division)

130. Consider the following actions that a DBMS can take in response to a constraint violation during an update operation.
a. Reject the update operation.
b. Perform the update and log the violation.
c. Trigger a user-defined error-correction routine.
d. Cascade the update to related tuples if a referential integrity constraint is involved.
Which of the above are valid actions?
(A) a and d only
(B) b and c only
(C) a, c, and d only
(D) a, b, c, and d

131. In an ER diagram, a multi-valued attribute is represented by:
(A) A dashed ellipse
(B) A rectangle
(C) An ellipse
(D) A double ellipse

132. Consider the following statements and select the correct option.
Statement I : Codd's rule of guaranteed access states that all data must be accessible without ambiguity.
Statement II : Codd's rules provide a benchmark for what constitutes a truly relational database system.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

133. Which SQL keyword is used to retrieve data from a database?
(A) `GET`
(B) `FETCH`
(C) `SELECT`
(D) `RETRIEVE`

134. A developer creates a named block of SQL statements that can be executed as a unit.
It can accept parameters but does not have to return a value.
This is a:
(A) View
(B) Trigger
(C) Function
(D) Stored Procedure

135. Consider the following potential remedies against SQL Injection attacks.
a. Using parameterized queries or prepared statements.
b. Validating and sanitizing all user input.
c. Using stored procedures.
d. Applying the principle of least privilege to database accounts.
e. Simply disabling all user input forms.
Which of the above statements are correct?
(A) a, b, and d only
(B) a, c, and e only
(C) a, b, c, and d only
(D) b, c, and e only

Downloaded from www.pdf.tube

Answer the question by selecting the appropriate option.

Assertion (A) : Data warehousing involves the process of constructing and using a data warehouse.

Reason (R) : A data warehouse is a subject-oriented, integrated, time-variant, and non-volatile collection of data in support of management's decision-making process.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

81. Which data mining task involves assigning items to one of several predefined categories or classes?
(A) Classification
(B) Clustering
(C) Regression
(D) Summarization

82. In data warehousing, what is a concept hierarchy used for?
(A) To define the relationships between tables in a star schema.
(B) To organize data into different levels of abstraction or detail.
(C) To control user access to the data warehouse.
(D) To schedule the loading of data into the warehouse.

83. The K-Nearest Neighbour (KNN) algorithm is an example of which type of data mining technique?
(A) Association Rule Mining
(B) Clustering
(C) Classification
(D) Regression

84. A company wants to group its customers into distinct segments based on their purchasing behavior, without any predefined labels for these segments.
Which data mining technique is most appropriate for this task?
(A) Classification
(B) Regression
(C) Association Rule Mining
(D) Clustering

85. Consider the following statements regarding characteristics of a Data Warehouse.
a. It is subject-oriented, focusing on topics like 'sales' or 'customers'.
b. It is integrated, combining data from multiple, heterogeneous sources.
c. It is time-variant, containing historical data over a long period.
d. It is non-volatile, meaning data is not updated or deleted in real-time but added as new information.
e. It is primarily designed for online transaction processing (OLTP).
Which of the above statements are correct?
(A) a, b, and c only
(B) c, d, and e only
(C) a, b, c, and d only
(D) a, b, c, d, and e

NODIA

86. Which of the following is NOT a common data mining technique?
(A) Classification
(B) Normalization
(C) Regression
(D) Clustering

87. What is the goal of the Regression technique in data mining?
(A) To group a set of objects in such a way that objects in the same group are similar.
(B) To find interesting relationships among variables in large datasets.
(C) To predict a continuous-valued output for a given input.
(D) To identify a compact description for a subset of data.

88. Consider the following statements and select the correct option.

Statement I : A Hidden Markov Model (HMM) is a statistical model used for modeling sequences of observable events.

Statement II : Support Vector Machine (SVM) is a supervised learning model used for classification and regression analysis.

(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.

Dependency

108. The three Vs that are commonly used to characterize Big Data are:
(A) Volume, Velocity, and Variety
(B) Volume, Velocity, and Variability
(C) Volume, Veracity, and Validity
(D) Variety, Vision, and Value

109. What is the core function of MapReduce in the Hadoop ecosystem?
(A) A distributed file system for storing large data sets.
(B) A NoSQL database for key-value data storage.
(C) A framework for processing large data sets in parallel across a cluster.
(D) A data warehousing solution for OLAP queries.

110. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : The Hadoop Distributed File System (HDFS) is designed to be highly fault-tolerant.
Reason (R) : HDFS achieves fault tolerance by replicating data blocks across multiple nodes in the cluster.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

111. Which of the following is NOT a type of Big Data?
(A) Structured Data
(B) Semi-structured Data
(C) Unstructured Data
(D) Normalized Data

112. In the context of Big Data architecture, what is the role of Hadoop?
(A) It is a relational database management system.
(B) It is an open-source framework for storing and processing Big Data.
(C) It is a data visualization tool.
(D) It is a cloud computing platform.

113. What does the Velocity characteristic of Big Data refer to?
(A) The speed at which data is generated and must be processed.
(B) The sheer size of the data.
(C) The different forms of data sources.
(D) The uncertainty or trustworthiness of the data.

114. The Map phase in a MapReduce job typically performs which actions?
(A) Aggregating and summarizing
(B) Filtering and sorting
(C) Storing and replicating
(D) Reducing and combining

115. What is the fundamental unit of data storage in HDFS?
(A) File
(B) Record
(C) Block
(D) Node

116. Which of the following best describes NoSQL databases?
(A) They serve as a direct replacement for all existing SQL and relational database management systems.
(B) They exclusively utilize standard SQL query syntax for the primary manipulation of all stored data.
(C) They are primarily used for online transaction processing (OLTP) while maintaining strict ACID guarantees.
(D) They provide storage and retrieval for data modeled by means other than tabular relations used in RDBMS.

117. Which of the following is an example of a NoSQL database product?
(A) MySQL
(B) MongoDB
(C) Oracle
(D) SQL Server

118. NoSQL databases are often described as Not Only SQL. What does this imply?
(A) They may support SQL-like query languages but are not based on the relational model.

Click Here to Buy book on Amazon

99. An Object-Relational Database supports all of the following except:
(A) User-defined data types
(B) Inheritance
(C) The pure relational model with no extensions
(D) Complex objects

100. Consider the following statements regarding query processing and optimization.
a. Query processing refers to the range of activities involved in extracting data from a database.
b. A key goal of query optimization is to minimize the execution time.
c. Cost-based optimization uses statistics about the data to estimate the cost of different execution plans.
d. Relational algebra equivalences are used to transform one query expression into a logically equivalent one.
Which of the above statements are correct?
(A) a and b only
(B) a, c, and d only
(C) b, c, and d only
(D) a, b, c, and d

101. Which of the following is NOT a concurrency control technique?
(A) Timestamp ordering
(B) Two-phase locking
(C) Validation technique
(D) Shadow paging

102. Consider the following statements and select the correct option.
Statement I : Normalization helps to eliminate insertion, updation, and deletion anomalies.
Statement II : A drawback of excessive normalization can be poorer query performance due to an increased number of joins.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

103. In database security, the `GRANT` command is used to:
(A) Remove privileges from a user.
(B) Create a new user account.
(C) Give privileges to a user.
(D) Encrypt the database content.

104. Which phenomenon occurs when a transaction T1 retrieves a set of rows, and a subsequent query by T1 within the same transaction retrieves a different set of rows because another transaction T2 has inserted new rows?
(A) Phantom Read
(B) Dirty Read
(C) Non-repeatable Read
(D) Lost Update

105. Consider the following statements regarding database systems.
a. A 'transaction' is a logical unit of work.
b. 'Concurrency control' is the management of simultaneous operations on the database without them interfering with one another.
c. 'Database recovery' is the process of restoring a database to a correct state after a failure.
d. 'Functional dependency' is a relationship that exists when one attribute uniquely determines another attribute.
e. 3NF is stricter than BCNF.
Which of the above statements are correct?
(A) a, b, and e only
(B) c, d, and e only
(C) a, b, c, and d only
(D) a, c, and d only

106. A relation R is in 3NF but not in BCNF. What can be concluded?
(A) R has a transitive dependency.
(B) R has a non-trivial functional dependency $X \rightarrow A$ where X is not a superkey, and A is a prime attribute.
(C) R has a partial dependency.
(D) R is not in 1NF.

107. Which of the following pairs is incorrectly matched?
(A) Database Recovery : Checkpointing
(B) Concurrency Control : Two-Phase Locking
(C) Database Security : GRANT/REVOKE
(D) Query Optimization : Functional

Downloaded from www.pdf.tube

(D) a, d, c, b

71. A Mobile Database system must be designed to handle which specific challenge?
(A) Terabytes of static data.
(B) High-speed network connectivity at all times.
(C) Complex logical deductions.
(D) Disconnected operation and data synchronization.

72. Consider the following statements and select the correct option.
Statement I : An XML database is a data persistence software system that allows data to be stored in XML format.
Statement II : Unlike relational databases, XML databases do not require a predefined schema.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

73. The client-server architecture where the client is primarily responsible for presentation and the server is responsible for data processing and storage is known as a:
(A) Fat client model
(B) Thin client model
(C) Two-tier architecture
(D) Three-tier architecture

74. What is a key feature of a Deductive Database?
(A) The ability to query data based on its content (e.g., image similarity).
(B) The ability to store and manage location-based data.
(C) The ability to track the history of data changes over time.
(D) The ability to infer new facts from existing data using rules.

75. Consider the following data models.
a. Multimedia Databases
b. Temporal Databases
c. Geographic Information Systems
d. Hierarchical Databases
Which of the above are considered enhanced data models, going beyond the traditional relational model?
(A) a and b only
(B) a, b, and c only
(C) c and d only
(D) a, b, c, and d

76. In OLTP systems, transactions are typically:
(A) Long, complex queries for analysis.
(B) Short, fast, and atomic operations.
(C) Read-only and involve large datasets.
(D) Executed in batch mode overnight.

77. Which data mining technique is used to find rules that predict the occurrence of an item based on the occurrences of other items in a transaction?
(A) Classification
(B) Clustering
(C) Regression
(D) Association Rule Mining

78. Which of the following is a key difference between OLTP and OLAP?
(A) OLTP is used for data analysis, while OLAP is used for transaction processing.
(B) OLTP databases are highly normalized, while OLAP databases (data warehouses) are often denormalized.
(C) OLTP queries are complex and involve aggregations, while OLAP queries are simple lookups.
(D) OLTP systems store historical data, while OLAP systems store only current data.

79. What does the 'Support' of an association rule X => Y measure?
(A) The percentage of transactions that contain Y among those that contain X.
(B) The likelihood that a transaction containing X will also contain Y.
(C) The percentage of transactions in the database that contain both X and Y.
(D) The ratio of the rule's support to the support of its consequent.

80. The following question consists of two statements — Assertion (A) and Reason (R).

(B) They completely reject SQL-like query languages.
(C) They are an older technology that predates SQL.
(D) They can only be used for storing non-relational data.

119. A key advantage of using NoSQL databases in a cloud environment is their:
(A) Strict consistency model.
(B) Vertical scalability.
(C) Horizontal scalability and flexibility.
(D) Reliance on predefined schemas.

120. Match the NoSQL database type in List I with its characteristic in List II and select the correct option.
List I: (NoSQL Type)
P. Key-Value Store
Q. Document Store
R. Column-Family Store
S. Graph Database
List II: (Characteristic)
1. Stores data in documents, often JSON or BSON.
2. Optimized for traversing relationships between entities.
3. Simplest model, data stored as a dictionary or hash table.
4. Stores data in tables with rows and dynamic columns.
(A) P-1, Q-3, R-4, S-2
(B) P-3, Q-1, R-4, S-2
(C) P-3, Q-4, R-1, S-2
(D) P-2, Q-1, R-4, S-3

121. In NoSQL systems, what is a primary concern for query optimization?
(A) Ensuring compliance with the BCNF normal form.
(B) Enforcing referential integrity across collections.
(C) Maximizing the use of foreign key constraints.
(D) Minimizing the number of joins, as they are often expensive or unsupported.

122. Consider the following statements and select the correct option.
Statement I : NoSQL databases typically offer weaker consistency guarantees than traditional relational databases to achieve higher availability and scalability.
Statement II : Indexing in NoSQL databases is crucial for performance, similar to relational databases, but the implementation can vary significantly based on the data model.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

123. A social media company wants to store user profiles and their connections.
Querying friends of friends should be very efficient.
Which type of NoSQL database would be most suitable?
(A) Key-Value Store
(B) Document Store
(C) Graph Database
(D) Column-Family Store

124. Which of the following is a common strategy for ordering data sets in NoSQL databases to improve query performance?
(A) Using foreign keys
(B) Normalizing the data to 3NF
(C) Storing all data in a single large table
(D) Using natural sort order of primary keys or clustering keys

125. Arrange the following data storage paradigms in the chronological order of their widespread adoption in the industry.
a. Relational Databases (SQL)
b. File Systems
c. NoSQL Databases
d. Object-Oriented Databases
Which of the following represents the correct sequence?
(A) b, a, d, c
(B) a, b, d, c
(C) b, d, a, c
(D) a, c, d, b

126. A user wants to define a new table in a database. Which category of database language would be

Click Here to Buy book on Amazon

(A) Link Analysis
(B) Dependency Modeling
(C) Sequencing Analysis
(D) Summarization

148. Which statement about Support Vector Machines (SVM) is most accurate?
(A) It is a supervised learning algorithm that finds an optimal hyperplane to separate classes.
(B) It is an unsupervised learning algorithm for clustering.
(C) It is a method for finding frequent itemsets in transactional data.
(D) It is a probabilistic model for sequential data.

149. In the context of Big Data, HDFS stands for:
(A) Hadoop Distributed File Storage
(B) High-Density File System
(C) Hadoop Distributed File System
(D) Hierarchical Data File System

150. Match the Big Data characteristic in List I with its description in List II and select the correct option.
List I: (Characteristic)
P. Volume
Q. Velocity
R. Variety
S. Veracity
List II: (Description)
1. The different types of data (structured, unstructured, etc.).
2. The trustworthiness or quality of the data.
3. The scale or amount of data.
4. The speed at which data is generated and processed.
(A) P-3, Q-4, R-1, S-2
(B) P-1, Q-2, R-3, S-4
(C) P-3, Q-1, R-4, S-2
(D) P-4, Q-3, R-2, S-1

151. The Reduce phase in a MapReduce job is responsible for:
(A) Distributing the data to different nodes.
(B) Processing each input record to generate intermediate key-value pairs.
(C) Aggregating the intermediate key-value pairs to produce the final output.
(D) Replicating the data blocks for fault tolerance.

152. Which NoSQL data model would be most appropriate for a content management system where articles, with varying attributes like author, title, body, and tags, are stored?
(A) Key-Value Store
(B) Document Store
(C) Column-Family Store
(D) Graph Database

153. A primary motivation for the development of NoSQL databases was the need for:
(A) Better support for SQL joins.
(B) Stricter ACID compliance.
(C) Greater horizontal scalability and handling of unstructured data.
(D) More robust normalization procedures.

154. Indexing in a NoSQL database primarily serves to:
(A) Enforce data normalization.
(B) Ensure ACID compliance.
(C) Provide data security.
(D) Improve the speed of data retrieval operations.

155. Which of the following represents a logical description of some portion of the database that is required by a particular user?
(A) External Schema or View
(B) Internal Schema
(C) Conceptual Schema
(D) Physical Schema

156. The ability to modify the conceptual schema without causing any changes in the external schemas or application programs is known as:
(A) Physical Data Independence
(B) Logical Data Independence
(C) Architectural Independence
(D) Concurrency Control

157. Which of the following is NOT a relational algebra operator?
(A) Union
(B) Difference
(C) Transitive Closure
(D) Cartesian Product

NODIA

136. What is the main characteristic of a functional dependency A → B?
(A) B depends on A.
(B) A determines B.
(C) A and B are both candidate keys.
(D) For a given value of B, there is only one possible value of A.

137. Which of the following is a cost component considered by a query optimizer?
(A) All of the above
(B) CPU time
(C) I/O operations
(D) Network communication cost

138. Which concurrency problem occurs when a transaction reads a value, another transaction updates it and commits, and then the first transaction re-reads it and gets a different value?
(A) Dirty Read
(B) Non-repeatable Read
(C) Lost Update
(D) Phantom Read

139. In a recovery system, what is the purpose of a 'checkpoint'?
(A) To mark the start of a transaction.
(B) To permanently delete old log records.
(C) To reduce the amount of log data that needs to be scanned during recovery.
(D) To create a physical backup of the entire database.

140. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : Mandatory Access Control (MAC) provides a higher level of security than Discretionary Access Control (DAC).
Reason (R) : MAC policies are controlled by a central system administrator and individual users cannot alter access control policies for the data they own, preventing mistakes and malicious behavior.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

141. A database tracking employee salaries over the past 20 years, allowing queries like What was John Smith's salary in 2010?, is an example of a:
(A) Deductive Database
(B) Mobile Database
(C) Multimedia Database
(D) Temporal Database

142. What does XML stand for in the context of internet databases?
(A) eXtended Markup Language
(B) eXtensible Markup Language
(C) eXecutable Markup Language
(D) eXternal Markup Language

143. Which of the following is NOT a feature of distributed databases?
(A) Location transparency
(B) Replication transparency
(C) Single point of failure
(D) Fragmentation transparency

144. Which data mining technique would be used to predict a student's final grade (a continuous value) based on their midterm scores and attendance?
(A) Classification
(B) Clustering
(C) Association Rule Mining
(D) Regression

145. In the star schema of a data warehouse, the central table is called the:
(A) Fact table
(B) Dimension table
(C) Attribute table
(D) Helper table

146. A data warehouse is described as non-volatile because:
(A) The data is never updated once it is added.
(B) The data is not subject to frequent changes and deletions as in an OLTP system.
(C) The data is stored on non-volatile storage like disks.
(D) The data is protected against volatility in the stock market.

147. Which data mining analysis method is used to discover patterns where the presence of one event is followed by another later in time?

Downloaded from www.pdf.tube

(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

89. Link analysis, used in data mining, is primarily concerned with:
(A) Predicting a numerical value.
(B) Grouping similar data points together.
(C) Discovering relationships and connections within data.
(D) Assigning data points to predefined classes.

90. Match the data mining technique in List I with its description in List II and select the correct option.
List I: (Technique)
P. Classification
Q. Clustering
R. Association Rule Mining
S. Regression
List II: (Description)
1. Predicts a continuous value.
2. Finds if-then patterns.
3. Groups data into segments based on similarity.
4. Assigns items to predefined categories.
(A) P-4, Q-3, R-2, S-1
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-2, R-3, S-1
(D) P-3, Q-4, R-1, S-2

91. Social Network Analysis is a data mining technique that focuses on:
(A) The summarization of large text documents.
(B) The analysis of sequences of events over time.
(C) The modeling of dependencies between different variables.
(D) The study of social structures through the use of network and graph theory.

92. A supermarket is analyzing its sales data and discovers the rule {Bread, Butter} => {Milk} with high confidence.
This is an example of:
(A) Sequencing Analysis
(B) Association Rule
(C) Link Analysis
(D) Summarization

93. Which technique would be most suitable for predicting the future price of a stock based on its past performance?
(A) Clustering
(B) Classification
(C) Regression
(D) Association Rules

94. What is the primary purpose of Dependency Modeling in data mining?
(A) To find which items are frequently bought together.
(B) To predict the class label of a new data instance.
(C) To analyze the evolution of data over time.
(D) To describe the significant dependencies among variables.

95. In the context of K-Nearest Neighbour (KNN), what does 'K' represent?
(A) The number of nearest neighbors to consider for classification.
(B) The number of clusters to be formed.
(C) The number of attributes in the dataset.
(D) The confidence threshold for an association rule.

NODIA

96. Which normal form is violated if a non-prime attribute is functionally dependent on a part of a composite candidate key?
(A) First Normal Form (1NF)
(B) Second Normal Form (2NF)
(C) Third Normal Form (3NF)
(D) Boyce-Codd Normal Form (BCNF)

97. The ACID property that ensures that once a transaction is committed, its effects are permanent, is:
(A) Durability
(B) Atomicity
(C) Consistency
(D) Isolation

98. A 'log record' in a database system is primarily used for:
(A) Query optimization
(B) Database recovery
(C) Concurrency control
(D) Security authorization

158. What is the effect of the `ON DELETE CASCADE` clause in a foreign key constraint?
(A) It prevents the deletion of a referenced row in the parent table.
(B) When a referenced row is deleted, the foreign key values in referencing rows are set to NULL.
(C) It raises an error and rolls back the transaction.
(D) When a referenced row is deleted, all corresponding referencing rows are also deleted.

159. Consider a relation R(Student_ID, Course_ID, Professor_Name) where a student can take multiple courses and a course can be taught by multiple professors.
The primary key is (Student_ID, Course_ID).
The functional dependency Professor_Name → Course_ID exists.
What is the most accurate description of the problem with this design?
(A) The functional dependency is incorrect and violates BCNF.
(B) It violates 1NF because Professor_Name is not atomic.
(C) It contains a partial dependency, violating 2NF.
(D) It contains a transitive dependency, violating 3NF.

160. Consider the following statements regarding advantages of using a DBMS.
a. Controlling redundancy
b. Providing security and authorization
c. Enforcing integrity constraints
d. Providing backup and recovery
Which of the above statements are correct?
(A) a and c only
(B) b and d only
(C) a, b, and d only
(D) a, b, c, and d

161. A user executes an `UPDATE` statement in SQL. This is an example of using:
(A) DDL
(B) DML
(C) DCL
(D) TCL

162. Consider the following statements and select the correct option.
Statement I : A view is a physical table that stores a subset of the database.
Statement II : Changes made to the data in a view can sometimes be reflected in the underlying base tables.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

163. The primary key of a relational table has two properties. They are:
(A) Sequential and Not Null
(B) Unique and Indexed
(C) Unique and Not Null
(D) Referenced and Not Null

164. A deadlock in a database system is a situation where:
(A) A transaction cannot proceed because the data it needs is corrupted.
(B) The system crashes due to a hardware failure.
(C) A query optimizer cannot find an efficient execution plan.
(D) Two or more transactions are waiting indefinitely for each other to release locks.

165. Consider the following statements regarding enhanced data models.
a. Temporal databases incorporate aspects of time.
b. Multimedia databases handle data like images, audio, and video.
c. Deductive databases use facts and rules to infer information.
d. Mobile databases are optimized for devices with limited power and intermittent connectivity.
e. Geographic Information Systems (GIS) manage spatial data.
Which of the above statements are correct?
(A) a, b, and e only
(B) c, d, and e only
(C) a, b, c, and d only

Click Here to Buy book on Amazon

245. Consider the following statements regarding SQL.

a. `DELETE` removes rows from a table, while `DROP` removes the entire table.

b. A `VIEW` can be used to simplify complex queries and restrict access to data.

c. A database `TRIGGER` is a set of SQL statements that run automatically when an event occurs.

d. `ALTER TABLE` is a DDL command used to modify the structure of an existing table.

Which of the above statements are correct?

(A) a and b only
(B) a, c, and d only
(C) b and d only
(D) a, b, c, and d

246. What is a key characteristic of a weak entity set in an ER diagram?

(A) It has its own primary key composed entirely of its own attributes.
(B) It does not have enough attributes to form a primary key on its own and depends on a strong entity.
(C) It can be uniquely identified without depending on another entity.
(D) It is represented by a single-lined rectangle.

247. In data mining, what is the main difference between classification and clustering?

(A) Classification works with numerical data, while clustering works with categorical data.
(B) Classification finds groups, while clustering predicts a value.
(C) Classification is used for sequencing, while clustering is used for association.
(D) Classification is supervised learning (predefined classes), while clustering is unsupervised learning (no predefined classes).

248. The architecture of a database system that includes an application server between the user's client and the database server is known as:

(A) One-tier architecture
(B) Three-tier architecture
(C) Two-tier architecture
(D) Centralized architecture

249. What is a primary characteristic of Big Data that makes it difficult to manage with traditional DBMS?

(A) Its massive volume, high velocity, and wide variety.
(B) Its reliance on the SQL query language.
(C) Its high degree of normalization.
(D) Its requirement for strict ACID compliance.

250. Consider the following statements regarding NoSQL databases.

a. They are designed for horizontal scalability.

b. They often have a flexible or non-existent schema.

c. They can handle structured, semi-structured, and unstructured data.

d. They are suitable for large-scale web applications and Big Data.

e. They are a single type of database with one specific data model.

Which of the above statements are correct?

(A) a, b, and c only
(B) c, d, and e only
(C) a, b, c, and d only
(D) a, b, and e only

NODIA

Direction : Please read the following source carefully.

The foundational architecture of a modern Database Management System (DBMS) is often conceptualized through the Three-Schema Architecture, which promotes data independence. This architecture delineates three levels of schemas: the internal schema, the conceptual schema, and the external schema. The internal schema describes the physical storage structure of the database, detailing data structures and file organizations. The conceptual schema provides a unified, logical representation of the entire database for the community of users, hiding the physical storage details. Finally, the external schema, or view level, describes the portion of the database that is relevant to a particular user group, concealing the rest of the database. This separation is crucial for achieving data independence. Logical data independence allows the conceptual schema to be changed without affecting external schemas, while physical data independence permits modifications to the in-ternal schema without altering the conceptual schema. Database languages,

(A) Relational table
(B) Tree structure
(C) Data Cube
(D) Graph structure

178. What is the purpose of the Confidence metric in association rule mining?
(A) It measures how frequently the items in the rule appear together across the entire available dataset.
(B) It measures the total number of individual transactions in the database that fully support the given rule.
(C) It measures the inherent complexity of the association rule based on the number of items it contains.
(D) It measures the reliability of the rule, indicating how often the given rule is found to be accurate.

179. The defining characteristic of unstructured data in the context of Big Data is that it:
(A) Has a predefined data model and is highly organized.
(B) Is always in text format.
(C) Does not conform to a traditional row-column database structure.
(D) Cannot be processed by computers.

180. Match the SQL term in List I with its purpose in List II and select the correct option.
List I: (SQL Term)
P. Trigger
Q. View
R. Stored Procedure
S. Constraint
List II: (Purpose)
1. A named block of reusable SQL code.
2. A rule to enforce data integrity.
3. A stored query that acts as a virtual table.
4. A procedure automatically executed in response to a DML event.
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-3, R-1, S-2
(C) P-4, Q-1, R-3, S-2
(D) P-2, Q-3, R-1, S-4

181. Which NoSQL database is modeled on Google's Bigtable and is organized as a sparse, distributed, persistent multi-dimensional sorted map?
(A) Redis
(B) Cassandra
(C) MongoDB
(D) Neo4j

NODIA

182. Consider the following statements and select the correct option.
Statement I : Data independence is the ability to modify a schema at one level of the database system without altering the schema at the next higher level.
Statement II : A client-server architecture is a form of distributed database system.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

183. In an ER diagram, attributes are represented by:
(A) Rectangles
(B) Diamonds
(C) Lines
(D) Ellipses

184. When dealing with constraint violations, setting a foreign key to a default value is an example of which remedial action?
(A) RESTRICT
(B) SET NULL / SET DEFAULT
(C) CASCADE
(D) TRIGGER

185. One of Codd's rules, the Comprehensive Data Sublanguage Rule, implies that:
(A) A relational system must support a language that can handle data definition, manipulation, and transaction management.
(B) The system must support null values.
(C) The database description must be stored at the logical level in the same way as ordinary data.
(D) Physical storage details must be hidden from the user.

186. Which SQL statement is used to delete data from a database table?
(A) `REMOVE FROM table_name WHERE condition;`

(B) `DELETE FROM table_name WHERE condition;`
(C) `DROP FROM table_name WHERE condition;`
(D) `TRUNCATE FROM table_name WHERE condition;`

187. Which of the following is NOT a property of a relation in the relational model?
(A) The tuples of the relation have no specific order.
(B) Each attribute value is atomic.
(C) The attributes of the relation have a specific order.
(D) Each tuple is distinct.

188. If a relation is in 3NF, it has already eliminated:
(A) Multi-valued dependencies only.
(B) Transitive dependencies only.
(C) All dependencies.
(D) Partial and transitive dependencies.

189. A transaction that has completed its execution successfully is said to be:
(A) Rolled back
(B) Aborted
(C) Committed
(D) Locked

190. Consider the following statements about query optimization.
a. It aims to find the most efficient way to execute a given query.
b. A query plan or execution plan is a sequence of relational algebra operations.
c. Heuristic-based optimization uses rules to transform a query into a more efficient form.
d. A query optimizer is an essential part of a modern RDBMS.
Which of the above statements are correct?
(A) a and d only
(B) b and c only
(C) a, c, and d only
(D) a, b, c, and d

191. The strict two-phase locking protocol (Strict 2PL) differs from the basic 2PL in that it:
(A) Releases all locks at the end of the transaction.
(B) Releases read locks as soon as the read is complete.
(C) Does not allow any locks to be released until the transaction commits or aborts.
(D) Allows lock conversion from write to read.

192. Which of the following is an example of Discretionary Access Control (DAC)?
(A) A system where data is labeled 'Top Secret' and can only be accessed by users with 'Top Secret' clearance.
(B) A system where the owner of a file can grant read/write access to other specific users.
(C) A system where access is determined by a user's role in an organization (e.g., 'Manager', 'Clerk').
(D) A system that automatically blocks all access during a security alert.

193. The concept of an object having its own private data and public methods to access that data is known as:
(A) Encapsulation
(B) Inheritance
(C) Polymorphism
(D) Abstraction

194. A developer needs to model complex biological data, including gene sequences and protein structures.
Which enhanced data model would be most appropriate?
(A) Geographic Information Systems
(B) Genome Data Management
(C) Temporal Database
(D) Deductive Database

195. Consider the following comparisons between OLAP and OLTP.
a. OLTP is transaction-oriented; OLAP is analysis-oriented.
b. OLTP systems deal with current data; OLAP systems deal with historical data.
c. OLTP databases are typically normalized; OLAP databases are typically denormalized.
d. OLTP queries are simple; OLAP queries are complex and involve aggregations.
e. OLTP is optimized for reads; OLAP is optimized for writes.
Which of the above statements are correct?
(A) a, b, and c only
(B) a, b, c, and d only

Downloaded from www.pdf.tube

(C) c, d, and e only
(D) a, b, c, d, and e

196. The Apriori algorithm is used for which data mining task?
(A) Classification
(B) Association Rule Mining
(C) Clustering
(D) Regression

197. A data mining analysis of website user logs reveals that 70% of users who visit page A, then visit page B, and finally make a purchase on page C.
This is an example of:
(A) Sequencing Analysis
(B) Summarization
(C) Dependency Modeling
(D) Link Analysis

198. The Veracity characteristic of Big Data refers to:
(A) The speed of data generation.
(B) The uncertainty and trustworthiness of data.
(C) The scale of the data.
(D) The different formats of data.

199. In the Hadoop ecosystem, the component responsible for scheduling jobs and managing cluster resources is:
(A) HDFS
(B) MapReduce
(C) YARN (Yet Another Resource Negotiator)
(D) HBase

200. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : Graph databases are highly effective for managing data with many complex relationships, such as social networks or recommendation engines.
Reason (R) : The graph data model, with its focus on nodes and edges, allows for fast traversal of relationships without the need for expensive join operations found in relational databases.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

201. Why might a developer choose a NoSQL database over a relational database for a new application?
(A) The application requires strict ACID transactions for all operations.
(B) The data schema is well-defined, stable, and not expected to change.
(C) The application's primary need is to perform complex joins across many tables.
(D) The application requires high scalability and a flexible data model for unstructured data.

202. Which of the following is a core component of the Big Data Architecture?
(A) Data Sources
(B) All of the above
(C) Data Storage
(D) Batch Processing / Real-time Processing

203. Which pair represents a procedural language and a non-procedural language, respectively?
(A) Relational Calculus, Relational Algebra
(B) SQL, QBE
(C) Relational Algebra, Relational Calculus
(D) DDL, DML

204. In an ER Diagram, a diamond shape represents:
(A) An entity set
(B) An attribute
(C) A primary key
(D) A relationship set

205. Arrange the following normal forms in the order of least restrictive to most restrictive.
a. Boyce-Codd Normal Form (BCNF)
b. First Normal Form (1NF)
c. Third Normal Form (3NF)
d. Second Normal Form (2NF)
Which of the following represents the correct sequence?
(A) b, d, c, a
(B) a, c, d, b
(C) b, c, d, a
(D) a, b, d, c

(B) c, d, and e only
(C) a, b, c, and d only
(D) a, b, c, d, and e

226. What is the primary purpose of Codd's 12 Rules?
(A) To define the syntax of the SQL language.
(B) To serve as a basis for evaluating whether a DBMS can be considered truly relational.
(C) To provide a step-by-step guide for database normalization.
(D) To list the requirements for achieving ACID properties.

227. An SQL injection attack typically works by:
(A) Flooding the database server with too many connections.
(B) Guessing user passwords.
(C) Exploiting buffer overflows in the database server software.
(D) Tricking the application into including malicious SQL code in a database query.

228. Which concurrency issue is prevented by the 'Isolation' property of ACID?
(A) System crashes causing data loss.
(B) Incomplete transactions being saved to the database.
(C) Interference between concurrently executing transactions.
(D) Violation of integrity constraints.

229. What type of database would be most suitable for storing and analyzing historical stock market data, with a focus on queries related to specific time ranges?
(A) Temporal Database
(B) Multimedia Database
(C) Deductive Database
(D) Mobile Database

230. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : The MapReduce programming model is inherently parallel.
Reason (R) : The 'Map' tasks can be executed independently on different subsets of the input data, and the 'Reduce' tasks can also work in parallel on different key groups.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

231. The main advantage of a document store NoSQL database like MongoDB is:
(A) Its flexible schema, allowing documents in the same collection to have different structures.
(B) Its ability to perform fast graph traversals.
(C) Its simple key-value data model.
(D) Its strict enforcement of referential integrity through foreign keys.

232. What is a key characteristic of the client-server architecture for a DBMS?
(A) The entire DBMS is installed on every user's machine.
(B) The processing load is distributed between a central server and multiple client machines.
(C) It only works with relational databases.
(D) It does not require a network connection.

233. The relational algebra operation used to combine tuples from two relations based on a common attribute is:
(A) Union
(B) Intersection
(C) Projection
(D) Join

234. In the context of Codd's rules, the Systematic Treatment of Null Values rule means:
(A) Null values are not allowed in the database.
(B) Null values must be treated consistently, primarily as unknown or inapplicable, and distinct from zero or empty strings.
(C) Null values should be automatically converted to a default value.
(D) Null values are only allowed in primary key columns.

235. To remove all rows from a table named `Sales` without deleting the table structure itself, which SQL command is generally faster than `DELETE FROM Sales;`?
(A) `TRUNCATE TABLE Sales;`
(B) `ERASE TABLE Sales;`
(C) `EMPTY TABLE Sales;`

Click Here to Buy book on Amazon

blocked.
(B) The corresponding foreign key values are set to NULL.
(C) The rows with the corresponding foreign key are deleted.
(D) The update on the referenced primary key also updates the corresponding foreign key values.

215. Which of the following SQL statements is a part of the Data Control Language (DCL)?
(A) `GRANT`
(B) `CREATE TABLE`
(C) `UPDATE`
(D) `COMMIT`

216. Which of the following is NOT a valid SQL data type?
(A) `INTEGER`
(B) `STRUCTURE`
(C) `DECIMAL(10, 2)`
(D) `BOOLEAN`

217. A database schema is said to be in 1NF if:
(A) There are no partial dependencies.
(B) There are no transitive dependencies.
(C) Every determinant is a candidate key.
(D) All attributes are single-valued.

218. Shadow paging is a technique used for:
(A) Database recovery
(B) Concurrency control
(C) Query optimization
(D) Database security

219. The process of analyzing data from different perspectives and summarizing it into useful information, often using multidimensional data models, is known as:
(A) OLTP (Online Transaction Processing)
(B) DBMS (Database Management System)
(C) OLAP (Online Analytical Processing)
(D) HDFS (Hadoop Distributed File System)

220. Consider the following data mining techniques.
a. K-Nearest Neighbour
b. Support Vector Machine
c. Relational Algebra
d. Hidden Markov Model
Which of the above are recognized as valid data mining techniques?
(A) a and c only
(B) a, b, and d only
(C) c and d only
(D) a, b, d, and e

NODIA

221. The MapReduce paradigm is composed of two main functions. What are they?
(A) Read and Write
(B) Store and Process
(C) Filter and Aggregate
(D) Map and Reduce

222. The term NOSQL primarily emphasizes:
(A) A strict rejection of any SQL-like syntax.
(B) A move away from the relational model's dominance for certain types of applications.
(C) Superiority over SQL in all use cases.
(D) A focus on numerical data processing only.

223. An administrator wants to give a user 'John' the ability to query the 'Employees' table.
Which command is appropriate?
(A) `GRANT SELECT ON Employees TO John;`
(B) `ALLOW SELECT ON Employees TO John;`
(C) `SET PERMISSION SELECT ON Employees FOR John;`
(D) `UPDATE John SET PERMISSION = SELECT ON Employees;`

224. In the context of the three-schema architecture, what does the conceptual schema represent?
(A) The physical storage details of the database.
(B) A high-level, implementation-independent view of the entire database.
(C) The views of individual user groups.
(D) The programming interface to the database.

225. Consider the following statements regarding relational database design and normalization.
a. The goal of normalization is to reduce data redundancy.
b. A relation is in 2NF if it is in 1NF and all non-prime attributes are fully functionally dependent on the primary key.
c. A relation is in 3NF if it is in 2NF and has no transitive dependencies.
d. BCNF is a stricter version of 3NF.
e. Denormalization is sometimes used to improve query performance.
Which of the above statements are correct?
(A) a, b, and e only

206. A database trigger is a procedure that is executed automatically in response to:
(A) A user logging into the system.
(B) A specific DML (INSERT, DELETE, UPDATE) event on a table.
(C) The database server starting up.
(D) A query that takes too long to execute.

207. A relation R(A, B, C, D) has the functional dependency {AB → C, C → D}.
The primary key is {A, B}. This relation contains a:
(A) Partial dependency
(B) Multi-valued dependency
(C) Transitive dependency
(D) No dependency anomalies

208. The main goal of concurrency control in a DBMS is to:
(A) Allow multiple transactions to execute simultaneously without causing inconsistencies.
(B) Ensure transactions are durable.
(C) Speed up the execution of a single transaction.
(D) Ensure the database schema is in BCNF.

209. An Object-Oriented Database Management System (OODBMS) integrates database capabilities with:
(A) Procedural programming languages.
(B) Functional programming paradigms.
(C) Object-oriented programming language capabilities.
(D) Logic-based rule systems.

210. Match the data mining technique in List I to its typical use case in List II and select the correct option.
List I: (Technique)
P. Clustering
Q. Classification
R. Regression
S. Association Mining
List II: (Use Case)
1. Predicting house prices based on features like area and location.
2. Identifying if an email is 'spam' or 'not spam'.
3. Finding which products are frequently purchased together in a supermarket.
4. Segmenting customers into different groups based on their purchasing habits.
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-3, R-2, S-1
(C) P-4, Q-2, R-1, S-3
(D) P-2, Q-4, R-1, S-3

211. What is a Data Model in the context of DBMS?
(A) The physical hardware on which the database is stored.
(B) A collection of conceptual tools for describing data, data relationships, and consistency constraints.
(C) A specific software product like Oracle or MySQL.
(D) The actual data stored in the database at a point in time.

212. Consider the following statements and select the correct option.
Statement I : In a two-tier architecture, the application logic can be on the client, the server, or split between them.
Statement II : A three-tier architecture introduces a middle layer, typically an application server, between the client and the database server.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

213. What is the Cartesian product (R × S) of two relations R and S?
(A) A relation containing all tuples that are in R but not in S.
(B) A relation containing all tuples that are in both R and S.
(C) A relation containing all possible pairs of tuples where the first tuple is from R and the second is from S.
(D) A relation containing all tuples that are in either R or S.

214. If a foreign key is defined to have the `ON UPDATE CASCADE` action, what happens?
(A) The update on the referenced primary key is

Downloaded from www.pdf.tube

(D) `CLEAR TABLE Sales;`

236. A functional dependency X → Y is trivial if:
(A) X is a subset of Y.
(B) Y is a subset of X.
(C) X and Y are disjoint.
(D) X is a superkey.

237. A timestamp-based protocol for concurrency control works by:
(A) Locking data items before use.
(B) Validating transactions only at the commit stage.
(C) Assigning a unique timestamp to each transaction and ordering operations based on these timestamps.
(D) Creating multiple versions of each data item.

238. Which type of enhanced data model would be used to build a system that can answer queries like Who are Sherlock Holmes's associates? by inferring relationships from a set of facts and rules?
(A) Temporal Database
(B) Multimedia Database
(C) Geographic Information System
(D) Deductive Database

239. The process of building a data model for a data warehouse often involves creating which two types of tables?
(A) Parent and Child tables
(B) Object and Relational tables
(C) Fact and Dimension tables
(D) Master and Transaction tables

240. Match the term in List I with its corresponding technology/concept in List II and select the correct option.
List I: (Term)
P. HDFS
Q. MapReduce
R. NoSQL
S. OLAP
List II: (Technology/Concept)
1. Data Cube
2. MongoDB
3. Distributed File System
4. Parallel Processing Framework
(A) P-3, Q-4, R-2, S-1
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-2, R-4, S-1
(D) P-1, Q-4, R-2, S-3

241. What is the role of a DataNode in HDFS?
(A) To store the file system metadata.
(B) To store and retrieve data blocks as instructed by the NameNode.
(C) To manage the cluster resources.
(D) To coordinate the execution of MapReduce jobs.

242. Consider the following statements and select the correct option.
Statement I : Key-value stores are the simplest form of NoSQL databases.
Statement II : Graph databases are specifically designed to handle data and relationships with a high degree of connectivity.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct, but Statement II is incorrect.
(D) Statement I is incorrect, but Statement II is correct.

NODIA

243. The I in ACID properties stands for Isolation. What is the primary guarantee of Isolation?
(A) All operations of a transaction are completed successfully.
(B) The database remains in a consistent state after the transaction.
(C) The effects of a transaction are not visible to other transactions until it is committed.
(D) The effects of a committed transaction are permanent.

244. A university database stores course information in a table with columns `CourseID`, `CourseName`, `InstructorID`, and `InstructorName`.
Functional dependencies include `CourseID` → `InstructorID` and `InstructorID` → `InstructorName`.
`CourseID` is the primary key.
Which normal form is violated by this design?
(A) 1NF
(B) 2NF
(C) BCNF
(D) 3NF

(D) a, b, c, d, and e

166. A data model for a data warehouse that consists of a central fact table and a set of associated dimension tables is called a:
(A) Snowflake Schema
(B) Star Schema
(C) Galaxy Schema
(D) Constellation Schema

167. A model that describes the microscopic states of a system based on its macroscopic observations, often used in speech recognition, is the:
(A) K-Nearest Neighbour model
(B) Support Vector Machine (SVM)
(C) Regression Model
(D) Hidden Markov Model (HMM)

168. Which of the following best describes the goal of Summarization as a data mining task?
(A) To produce a compact and concise description of a dataset.
(B) To predict future values.
(C) To group similar items.
(D) To find if-then rules.

169. What is the role of a NameNode in HDFS?
(A) It stores the actual data blocks.
(B) It executes the Map tasks.
(C) It manages the file system namespace and metadata for all files and directories.
(D) It executes the Reduce tasks.

170. The following question consists of two statements — Assertion (A) and Reason (R). Answer the question by selecting the appropriate option.
Assertion (A) : NoSQL databases are often a good choice for applications with large amounts of unstructured data and high throughput requirements.
Reason (R) : Many NoSQL systems sacrifice strong consistency (as defined by ACID) in favor of performance, scalability, and availability (as described by the BASE properties).
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

171. A query language that specifies what data is to be retrieved without specifying how to retrieve it is called:
(A) Procedural
(B) Navigational
(C) Non-procedural or Declarative
(D) Imperative

172. A relation is in BCNF if and only if every determinant is a:
(A) Prime attribute
(B) Non-prime attribute
(C) Foreign key
(D) Superkey

173. The 'A' in the ACID properties of a transaction stands for:
(A) Atomicity
(B) Association
(C) Aggregation
(D) Authorization

174. Which of the following is a key component of a Geographic Information System (GIS)?
(A) Temporal data types
(B) Spatial data representation and querying
(C) Logical inference engine
(D) Support for disconnected clients

175. Identify the correct logical sequence of the following database design phases.
a. Implementation and loading
b. Physical design
c. Conceptual design
d. Logical design (data model mapping)
Which of the following represents the correct sequence?
(A) c, d, b, a
(B) d, c, a, b
(C) c, b, d, a
(D) b, c, d, a

176. Which data mining task is most similar to unsupervised learning?
(A) Classification
(B) Clustering
(C) Regression
(D) Association Rule Mining

177. The conceptual model for OLAP data is often represented as a:

Click Here to Buy book on Amazon

such as Data Definition Language (DDL) and Data Manipulation Language (DML), serve as interfaces to these schemas. Architecturally, DBMS can be deployed as centralized systems, where all components reside on a single computer, or as client/server architectures, where functionality is split between server processes and client applications, facilitating distributed access.

Based on the above source, attempt the following five questions.

251. Which component of the Three-Schema Architecture provides a unified, logical view of the entire database for all users?

A The external schema, which is tailored for specific user groups.
B The conceptual schema, which hides physical implementation details.
C The internal schema, which details the physical storage.
D The client/server schema, which defines distributed access.

252. The ability to modify the physical storage structure of a database without requiring changes to the application programs that use it is known as:

A Physical data independence.
B Logical data independence.
C Architectural data independence.
D Conceptual data independence.

253. According to the provided text, what is the primary purpose of an external schema?

A To define the physical data structures and file organization.
B To provide a comprehensive logical model of the entire database.
C To present a customized and relevant portion of the database to a user group.
D To split functionality between client applications and server processes.

254. Based on the description, one can infer that the separation of schemas in the Three-Schema Architecture primarily serves to:

A increase the complexity of database languages like DDL and DML.
B facilitate flexibility and evolution of the database system over time.
C enforce a centralized architecture for better performance.
D facilitate flexibility and evolution of the database system over time.

255. Which of the following statements accurately contrasts the roles of database languages as described in the passage?

A DDL is used to manipulate data within the external schema, while DML defines the internal schema.
B Both DDL and DML are used exclusively in client/server architectures, not centralized ones.
C These languages function as interfaces to define and manipulate the structures described by the schemas.
D DML is for creating views, while DDL is for managing physical storage directly.

Direction : Please read the following source carefully.

The relational model stands as a cornerstone of modern database design, organizing data into tables of rows and columns. Its formal foundation is supported by specific integrity constraints that ensure data accuracy and consistency. The two primary constraints are entity integrity and referential integrity. Entity integrity mandates that every table must have a primary key, and the column or columns chosen to be the primary key must be unique and not null. This ensures that every row in a table can be uniquely identified. Referential integrity, on the other hand, concerns the relationships between tables. It is maintained using a foreign key, which is a column or a set of columns in one table that refers to the primary key of another table. This constraint ensures that a value in the foreign key column either matches a primary key value in the referenced table or is null. The relational model is also equipped with formal query languages: Relational Algebra and Relational Calculus. Relational Algebra is a procedural language specifying the sequence of operations to obtain a result, while Relational Calculus is a non-procedural language describing the desired result without specifying how to compute it. Codd's rules provide a set of thirteen guidelines that a DBMS must satisfy to be considered truly relational.

Based on the above source, attempt the following five questions.

Downloaded from www.pdf.tube

256. What is the principal requirement of the entity integrity constraint in the relational model?
- A That a foreign key must always match a primary key in another table.
- B That every row in a table must have a unique, non-null primary key.
- C That all columns in a table must contain atomic values.
- D That data must be organized into a non-procedural structure.

257. A database designer ensures that an `Employee` table's `DepartmentID` column can only contain values that already exist in the `Department` table's `DepartmentID` primary key. This is an application of:
- A entity integrity.
- B Codd's rules.
- C relational calculus.
- D referential integrity.

258. Which of the following best distinguishes Relational Algebra from Relational Calculus?
- A Relational Algebra is procedural, defining a sequence of operations, while Relational Calculus is non-procedural.
- B Relational Algebra is non-procedural, whereas Relational Calculus is procedural.
- C Relational Algebra describes what result is desired, while Relational Calculus specifies the steps to get it.
- D Relational Algebra is used for defining constraints, while Relational Calculus is used for updating operations.

259. Based on the passage, the function of a foreign key is to:
- A guarantee that every row in its own table is unique.
- B act as the primary means of enforcing entity integrity.
- C establish and maintain a logical link between two tables.
- D provide a non-procedural method for querying data.

260. According to the text, Codd's rules serve as:
- A a benchmark for evaluating the relational fidelity of a DBMS.
- B a procedural language for data manipulation.
- C a set of constraints for ensuring data accuracy.
- D a non-procedural method for designing database schemas.

Direction : Please read the following source carefully.

SQL, or Structured Query Language, is the standard language for interacting with relational databases, providing a comprehensive suite of commands for data management. These commands are broadly categorized into Data Definition Language (DDL) and Data Manipulation Language (DML). DDL statements, such as `CREATE`, `ALTER`, and `DROP`, are used to define and manage the database schema, including tables and their data types. Constraints like `PRIMARY KEY`, `FOREIGN KEY`, `NOT NULL`, and `UNIQUE` are also specified using DDL to enforce data integrity. DML statements, including `INSERT`, `DELETE`, and `UPDATE`, are used for managing the data within the tables. Beyond basic data handling, SQL supports more advanced server-side objects to encapsulate logic. Views are virtual tables based on the result-set of an SQL statement, which can simplify complex queries and enhance security by restricting data access. Stored procedures and functions are pre-compiled collections of one or more SQL statements stored in the database, which can be executed by name, improving performance and promoting code reuse. Database triggers are a special type of stored procedure that automatically executes in response to specific DML events (`INSERT`, `UPDATE`, `DELETE`) on a table. However, the dynamic nature of SQL queries also exposes a significant security vulnerability known as SQL injection, where malicious code is inserted into input strings to be executed by the database.

Based on the above source, attempt the following five questions.

261. Which of the following SQL commands would be classified as a Data Definition Language (DDL) statement?
- A An `UPDATE` statement modifying a customer's address.
- B A `DELETE` statement removing an obsolete record.
- C An `INSERT` statement adding a new product to inventory.
- D An `ALTER` statement adding a new column

to a table.

262. A database administrator creates a virtual table named `High_Value_Customers` that only shows customers with purchases over a certain amount. According to the text, this object is best described as a:
A Stored Procedure.
B View.
C Database Trigger.
D Constraint.

263. What is the defining characteristic of a database trigger as described in the passage?
A It is a pre-compiled collection of SQL statements that must be called by name.
B It is a virtual table used to simplify complex query logic.
C It is a procedure that executes automatically in response to a DML event.
D It is a security mechanism to prevent SQL injection attacks.

264. The passage suggests that a primary benefit of using stored procedures and functions is to:
A define the data types and constraints for new tables.
B prevent malicious code from being inserted into input strings.
C provide a virtual representation of a subset of data.
D improve performance and encourage the reuse of SQL code.

265. SQL injection is presented as a security vulnerability that works by:
A inserting malicious SQL code into data inputs for execution.
B triggering an automatic procedure to delete sensitive data.
C altering the database schema using unauthorized DDL commands.
D creating views that expose confidential information.

Direction : Please read the following source carefully.

Normalization is a systematic process in relational database design aimed at minimizing data redundancy and improving data integrity. The process is guided by the analysis of functional dependencies, where the value of one attribute in a relation determines the value of another attribute. By organizing columns and tables, normalization helps to avoid undesirable characteristics such as insertion, update, and deletion anomalies. The output of the normalization process is a set of relations in a specific normal form (e.g., 1NF, 2NF, 3NF, BCNF). A well-normalized database design can significantly impact query performance. Efficient algorithms for query processing and optimization are employed by the DBMS to execute user queries in the most efficient way, and a normalized structure often provides the optimizer with better-defined paths for data retrieval. Beyond structural design, the operational aspect of a database is managed through transaction processing. A transaction is a logical unit of work that must be executed in its entirety or not at all. Concurrency control techniques are essential to manage simultaneous transactions and prevent them from interfering with each other, thus maintaining data consistency. In case of system failures, database recovery techniques, such as logging and checkpointing, are employed to restore the database to a consistent state, ensuring durability and atomicity.

Based on the above source, attempt the following five questions.

266. According to the text, the primary motivation for performing normalization on a relational database is to:
A increase the speed of query processing algorithms directly.
B reduce data redundancy and prevent modification anomalies.
C manage the simultaneous execution of transactions.
D restore the database to a consistent state after a failure.

267. The concept that is used as the fundamental guide for the process of normalization is:
A transaction processing.
B functional dependencies.
C concurrency control techniques.
D database recovery techniques.

268. A database system that ensures multiple users can access and modify data concurrently without corrupting it is employing:

Click Here to Buy book on Amazon

distributed, scale-out architecture aligns perfectly with the elastic and on-demand nature of cloud infrastructure, making them a popular choice for modern application development.

Based on the above source, attempt the following five questions.

291. The primary motivation behind the development of NOSQL databases was to address the need for:
- A strict data consistency and ACID transactions.
- B a standardized query language across all database types.
- C scalability and flexibility for web-scale applications.
- D a rigid schema to enforce data integrity.

292. According to the passage, how is data typically queried in a NOSQL database?
- A Using specific APIs or query languages tailored to the data model.
- B Exclusively through the standard Structured Query Language (SQL).
- C Primarily through complex join operations across multiple tables.
- D By relying on a predefined and rigid database schema.

293. The text suggests that query optimization in NOSQL systems is challenging primarily because of the:
- A overuse of the standard SQL language.
- B strict consistency requirements.
- C inherent unsuitability for cloud environments.
- D lack of a rigid schema and complex joins.

294. What is identified as a crucial factor for achieving good query performance in NOSQL databases?
- A The use of a relational data model.
- B The effective indexing and ordering of data sets.
- C The enforcement of strict ACID properties.
- D The normalization of the data to third normal form.

295. The passage highlights a strong alignment between NOSQL databases and cloud computing because of their:
- A support for a standard SQL interface.
- B rigid and structured schema design.
- C distributed, scale-out architecture.
- D emphasis on single-server performance.

Direction : Please read the following source carefully.

Ensuring the security and integrity of data is a critical function of any database management system. Database security involves protecting the database from unauthorized access, malicious attacks, and accidental damage. A core mechanism for implementing security is authorization, which is the process of defining and granting access rights or privileges to users and applications. These privileges dictate what data a user can see and what actions (e.g., SELECT, INSERT, UPDATE, DELETE) they can perform. In a distributed database environment, these challenges are magnified. A distributed database is a system where data is stored across several physical locations, and the DBMS manages the entire collection as a single logical database. The client-server architecture is commonly used to implement such systems, where client applications request data from one or more server processes that may be geographically dispersed. In this context, robust security and authorization protocols are paramount. The system must not only authenticate users but also enforce access controls consistently across the network, ensuring that a user's privileges are correctly applied regardless of which data server is processing their request. This prevents unauthorized access to sensitive data segments that may be stored on different nodes of the distributed system.

Based on the above source, attempt the following five questions.

NODIA

296. What is the primary function of authorization in the context of database security?
- A To protect the database from accidental hardware failures.
- B To store data across several different physical locations.
- C To ensure the database is logically perceived as a single entity.
- D To define and grant access rights and privileges to users.

297. According to the passage, why are security challenges magnified in a distributed database system?

C A GIS to manage the location of users.
D A multimedia database to store user avatars.

275. The challenge of querying data based on its content, such as finding all images that contain a specific object, is a primary concern for which database model?
A Deductive Databases.
B Temporal Databases.
C Multimedia Databases.
D Mobile Databases.

Direction : Please read the following source carefully.

Data warehousing represents a paradigm shift from transactional data processing to analytical data processing. A data warehouse is defined as a subject-oriented, integrated, time-variant, and non-volatile collection of data in support of management's decision-making process. Unlike operational databases that handle Online Transaction Processing (OLTP), data warehouses are optimized for Online Analytical Processing (OLAP). OLTP systems are characterized by frequent, short transactions involving updates, insertions, and deletions, while OLAP systems involve complex queries on large volumes of historical data. The data modeling approach for data warehouses also differs, often employing dimensional models like star or snowflake schemas. A crucial component of this analytical environment is the concept hierarchy, which allows data to be viewed and analyzed at different levels of abstraction. For example, a location hierarchy might range from individual stores to cities, states, and countries. This structure enables powerful OLAP operations such as roll-up, which involves aggregating data to a higher level of abstraction (e.g., from city sales to country sales), and drill-down, which navigates from less detailed data to more detailed data (e.g., from quarterly sales to monthly sales). These capabilities are fundamental for exploratory data analysis and business intelligence.

Based on the above source, attempt the following five questions.

276. Which of the following is identified as a defining characteristic of a data warehouse?
A It is a non-volatile and time-variant collection of data.
B It is optimized for frequent, short update transactions.
C It contains highly volatile and constantly changing data.
D It is designed for operational tasks rather than decision support.

277. The primary difference between OLTP and OLAP systems, as described in the passage, is that:
A OLTP focuses on complex analytical queries, while OLAP handles short transactions.
B OLTP uses historical data, while OLAP uses current operational data.
C OLTP supports management's decision-making, while OLAP supports daily business processes.
D OLTP involves frequent data modifications, while OLAP involves complex queries on historical data.

278. An analyst who aggregates daily sales data to view monthly sales trends is performing which type of OLAP operation?
A Drill-down.
B Roll-up.
C Slicing.
D Pivoting.

279. What is the main function of a concept hierarchy in the context of a data warehouse?
A To enforce data integrity constraints in transactional systems.
B To integrate data from various heterogeneous sources.
C To allow data analysis at different levels of abstraction.
D To model data using a star or snowflake schema.

280. The term subject-oriented in the definition of a data warehouse implies that the data is organized around:
A specific business processes like order entry.
B chronological time periods for historical analysis.
C different levels of data granularity.
D major subjects of the enterprise, such as customers or products.

Direction : Please read the following source

Downloaded from www.pdf.tube

ANSWERS

1. (B) The overall design or description of the database structure.
A database schema represents the logical configuration of an entire relational database. It serves as a blueprint that defines how data is organized and how the relationships among them are formulated. This description includes the layout of tables, types of data, and constraints, remaining relatively static over time compared to the actual data stored in the database at any moment.

2. (A) Both Statement I and Statement II are correct.
Centralized database architectures involve a single computer system where both the management software and the actual data reside for processing. In contrast, client-server architectures distribute tasks between a server providing database functionality and clients that request services. Both statements accurately reflect how these architectures handle the physical and logical distribution of database system components across various computing environments.

3. (C) The Entity-Relationship Model
A data model provides a set of conceptual tools used to describe data, relationships, and constraints within a database. The Entity-Relationship model is a prominent example that uses entities and their associations to represent complex data structures. It offers a high-level abstraction for database design, distinct from specific database software products or the low-level programming implementation details of the system.

4. (A) P-3, Q-4, R-1, S-2
Internal schemas define physical storage structures, while conceptual schemas describe the overall structure for the entire user community. External schemas focus on specific portions of the database relevant to particular user groups. Data independence allows modifications at one schema level without impacting the next higher level, ensuring that changes to storage or logic are effectively isolated within the three-schema architecture.

5. (C) Instance
A database instance refers to the actual collection of data stored within a database at a specific moment in time. While the schema defines the structure and design, the instance represents the current snapshot of information. This state changes frequently as users perform operations such as inserting, updating, or deleting records, reflecting the dynamic content stored within the predefined database schema.

6. (B) Data Processing Language (DPL)
Standard database systems utilize Data Definition Language for structure, Data Manipulation Language for handling data, and Data Control Language for security management. Data Processing Language is not recognized as a standard category of database language. These specialized languages provide the necessary interface for users and administrators to interact with the database management system and manage stored information according to specific needs.

NODIA

7. (D) Physical Data Independence
Physical data independence refers to the ability to modify the internal schema without requiring changes to the conceptual schema. This allows administrators to optimize physical storage structures or change hardware configurations without affecting the logical design of the database. It is a key feature of the three-schema architecture, ensuring that physical implementation details remain hidden from users and application programs.

8. (A) Both A and R are true and R is the correct explanation of A.
The three-schema architecture is specifically designed to provide data independence by separating user views from physical storage details. This structure divides the database into external, conceptual, and internal levels, allowing changes at one level without affecting others. The

fundamentally designed for scalability and fault tolerance. A prominent example of this is the Hadoop ecosystem. At its heart, Hadoop consists of two main components: the Map-Reduce programming model and the Hadoop Distributed File System (HDFS). Map-Reduce is a paradigm for processing and generating large data sets with a parallel, distributed algorithm on a cluster. It consists of a Map job, which processes and filters the data, and a Reduce job, which performs a summary or aggregation. HDFS is the storage component, a distributed file system designed to run on commodity hardware. It provides high-throughput access to application data and is highly fault-tolerant by replicating data blocks across multiple nodes in the cluster. This architecture allows organizations to store and process petabytes of data in a cost-effective and reliable manner, forming the backbone of many Big Data applications.

Based on the above source, attempt the following five questions.

286. Which of the following is identified in the text as a core characteristic of Big Data?
- A Consistency, referring to the uniform structure of the data.
- B Simplicity, indicating ease of management.
- C Veracity, indicating the trustworthiness of the data.
- D Velocity, referring to the high speed of data generation.

287. The Map-Reduce programming model is described as a paradigm for:
- A storing large datasets in a single, centralized repository.
- B ensuring the security of data stored in the Hadoop ecosystem.
- C processing large datasets in parallel across a distributed cluster.
- D providing a query language similar to SQL for Big Data.

288. According to the passage, what is the primary role of the Hadoop Distributed File System (HDFS)?
- A To execute the Map and Reduce jobs for data processing.
- B To provide a reliable and distributed storage layer for Big Data.
- C To manage the variety of structured and unstructured data types.
- D To handle the high velocity of incoming data streams.

289. The text states that HDFS achieves fault tolerance primarily through which mechanism?
- A By using high-end, specialized server hardware.
- B By encrypting all data blocks before storage.
- C By running a single, highly available master node.
- D By replicating data blocks across multiple nodes.

290. Based on the description, the function of the Reduce job in the Map-Reduce model is to:
- A perform aggregation or summarization of the mapped data.
- B filter and process the initial input data.
- C distribute the data across different nodes in the cluster.
- D replicate data for fault tolerance.

Direction : Please read the following source carefully.

NOSQL, which stands for Not Only SQL, represents a broad class of database management systems that differ from the classic relational model. These systems were developed to meet the modern challenges of web-scale applications, prioritizing scalability, performance, and flexibility over the strict consistency and structured schema of RDBMS. There are different NOSQL products, each with a unique data model, such as key-value stores, document databases, column-family stores, and graph databases. Querying and managing data in these systems often deviate from the standard SQL syntax, utilizing specific APIs or query languages tailored to their respective data models. For example, a document database might be queried using a JSON-based query structure. A key challenge in NOSQL systems is query optimization, which must be approached differently due to the absence of a rigid schema and complex joins. Performance is heavily reliant on effective indexing and ordering of data sets. Without proper indexing, queries can become very inefficient, especially in large-scale deployments. NOSQL databases are particularly well-suited for cloud computing environments, as their

Click Here to Buy book on Amazon

A concurrency control techniques.
B normalization algorithms.
C query optimization.
D database recovery techniques.

269. Based on the passage, what is the relationship between normalization and query optimization?
A Normalization replaces the need for query optimization algorithms.
B Query optimization is a technique used to achieve higher normal forms.
C A normalized schema can provide clearer data retrieval paths for the query optimizer.
D Normalization and query optimization are conflicting goals in database design.

270. A transaction is described as a logical unit of work. This implies that its operations are:
A independent and can be executed in any order.
B optimized for the fastest possible query performance.
C dependent on functional dependencies for their structure.
D indivisible and must be completed as a whole.

Direction : Please read the following source carefully.

Beyond the traditional relational model, enhanced data models have been developed to address the requirements of specialized applications. Temporal databases introduce concepts for managing time-varying data, systematically handling information that has a valid time or transaction time component, which is crucial for historical data analysis. Multimedia databases are designed to store, retrieve, and query large collections of multimedia objects like images, audio, and video, which are typically unstructured and require content-based retrieval mechanisms. Deductive databases extend the relational model by incorporating logic programming capabilities, allowing users to define rules and make logical inferences from the stored facts. For handling semi-structured data, which does not conform to a rigid schema, XML and Internet databases provide a flexible format and query languages like XQuery. The proliferation of smart devices has led to the development of mobile databases, which are lightweight databases that can operate on devices with limited resources and often in disconnected environments, requiring synchronization capabilities. Geographic Information Systems (GIS) are specialized database systems for capturing, storing, analyzing, and managing geospatial data, supporting complex spatial queries. Furthermore, the field of bioinformatics relies on genome data management systems to handle the vast and complex biological sequence data generated by research. These models showcase the evolution of databases to manage diverse and complex data types.

Based on the above source, attempt the following five questions.

271. A database system designed specifically to handle historical records and track changes to data over time would be best classified as a:
A Mobile Database.
B Multimedia Database.
C Temporal Database.
D Deductive Database.

272. Which type of enhanced data model is explicitly mentioned as being suitable for managing data in disconnected environments with limited resources?
A XML and Internet Databases.
B Mobile Databases.
C Geographic Information Systems.
D Genome Data Management Systems.

273. The key feature that distinguishes a deductive database from a standard relational database is its ability to:
A store and query large collections of images and videos.
B manage geospatial and spatial query types.
C handle semi-structured data without a rigid schema.
D use logical rules to infer new facts from existing data.

274. A web application developer needs to store and query data that is semi-structured, with nested elements and no fixed schema. Based on the text, which database model would be the most appropriate choice?
A An XML database for its flexible format.
B A temporal database to track the data's validity.

NODIA

A Because data is stored across multiple physical locations.
B Because data is stored in a single logical database.
C Because client applications make requests to servers.
D Because users are granted SELECT, INSERT, and UPDATE privileges.

298. A user successfully logs into a distributed banking system. The system then prevents this user from viewing account details from a different branch. This enforcement of access rights is an example of:
A database distribution.
B authorization.
C client-server communication.
D data replication.

299. The passage describes a distributed database as a system where:
A multiple client applications connect to a single, centralized server.
B data is stored in different locations but is managed as a single logical database.
C security is handled exclusively by the client applications.
D authorization is not required due to the physical separation of data.

300. Based on the text, in a distributed client-server architecture, a robust security protocol must ensure that:
A user privileges are enforced consistently across all data nodes.
B all data is stored on a single, highly secure server.
C client applications are responsible for their own security.
D every user has the same level of access to all data.

www.nodia.in

NODIA

carefully.

Data mining is the process of discovering patterns and valuable information from large datasets. It employs a variety of techniques to achieve different analytical goals. Association rule mining is a popular method used to discover interesting relationships between variables in large databases; a classic example is market basket analysis, which identifies products frequently purchased together. For predictive tasks, classification is used to assign items to predefined categories or classes. A model is first built from a training set of labeled data and then used to predict the class of new, unlabeled data. In contrast, clustering is an unsupervised technique that groups a set of objects in such a way that objects in the same group (or cluster) are more similar to each other than to those in other groups. This is useful when the class labels are unknown. Regression analysis is another predictive technique, but it is used to predict continuous values rather than discrete classes. Several specific algorithms are employed to implement these techniques. For instance, the Support Vector Machine (SVM) is a powerful supervised learning model used for both classification and regression. The K-Nearest Neighbour (KNN) algorithm is a simple, instance-based learning method where an object is classified by a majority vote of its neighbors. Other advanced techniques include Hidden Markov Models for sequential data, and social network analysis for modeling relationships.

Based on the above source, attempt the following five questions.

281. A retail company analyzes its sales data to find that customers who buy bread also tend to buy milk. Which data mining technique is being used?

A Association rule mining.
B Classification.
C Clustering.
D Regression.

282. What is the fundamental difference between classification and clustering as described in the text?

A Classification predicts continuous values, while clustering predicts categorical labels.
B Classification is a supervised technique using labeled data, while clustering is unsupervised.
C Classification groups similar objects, while clustering discovers relationships between variables.
D Classification uses the K-Nearest Neighbour algorithm, while clustering uses Support Vector Machines.

283. A data scientist is tasked with building a model to predict the exact selling price of a house based on its features like size and location. Which technique would be most appropriate for this task?

A Clustering.
B Regression.
C Association rule mining.
D Classification.

284. The K-Nearest Neighbour (KNN) algorithm is described as an instance-based method that classifies an object based on:

A the class that a majority of its neighbors belong to.
B a predefined set of logical rules.
C its distance from a calculated cluster center.
D a separating hyperplane that maximizes the margin.

285. A bank wants to segment its customers into distinct groups based on their banking habits without any predefined notion of what these groups might be. This task is a direct application of:

A classification, as it assigns customers to categories.
B regression, as it predicts customer behavior.
C clustering, as it groups similar customers without prior labels.
D association, as it links different customer actions.

Direction : Please read the following source carefully.

The emergence of Big Data is defined by its core characteristics, often summarized by several 'V's, including Volume (enormous scale of data), Velocity (high speed of data generation and processing), and Variety (diverse forms of data, both structured and unstructured). Handling such data requires a departure from traditional database architectures. Big Data architecture is

Downloaded from www.pdf.tube

reason accurately explains how this separation ensures that the physical representation of data remains independent from the logical view provided to users.

9. (C) To define the database schema and its structure.
The primary role of a Data Definition Language is to describe and define the database schema and its overall structure. It allows users to create, alter, or drop database objects like tables, indexes, and constraints. By using specific DDL commands, administrators establish the logical framework that dictates how data is organized, stored, and related within the database management system.

10. (A) Relationships between tables are represented by pointers.
Relational models represent data as a collection of tables where each row is a tuple and each column is an attribute. Unlike older hierarchical or network models, relationships in a relational database are represented by shared values across tables rather than physical pointers. Integrity constraints are also utilized within this model to maintain the accuracy and consistency of the stored information.

11. (D) a, b, c, and d
Codd's rules provide a standardized benchmark for relational database management systems. Rule 1 states all information must be in tables, while Rule 2 ensures accessibility via table names and keys. Rule 3 covers systematic null value treatment, and Rule 6 requires updatable views to be handled. All these statements are correct as they represent fundamental requirements for relational database systems.

12. (B) A weak entity set
In Entity-Relationship diagrams, a weak entity set is represented by a double-lined rectangle. These entities do not possess a sufficient number of attributes to form a primary key on their own and instead depend on a strong identifying entity. This visual notation helps designers distinguish entities that cannot exist independently from those that have their own unique and primary identification.

13. (C) Relation : Table
In the relational model, a relation is synonymous with a table, representing the fundamental unit for storing data. A tuple corresponds to a row within that table, and an attribute represents a column. These terms provide the formal mathematical foundation for relational databases, ensuring information is organized consistently through uniquely named tables, their constituent rows, and their defined attribute columns.

14. (B) b, a, c, d
Developing a database involves several chronological phases. First, requirements are collected and analyzed to understand user needs. This is followed by conceptual design to create a high-level model. Next, logical design maps this model to a specific database type, and finally, physical design determines storage details. This structured sequence ensures a logical progression from initial requirements to actual database implementation.

NODIA

15. (D) a, b, c, d, and e
Relational algebra is a procedural language specifying how to obtain results through operations, whereas relational calculus is a non-procedural language describing what results are needed. Both paradigms are formally equivalent in expressive power, providing different ways to query a database. They form the theoretical basis for SQL, allowing for the precise manipulation and retrieval of data in relational systems.

16. (B) To establish and enforce a link between data in two tables.
Foreign key constraints are used to establish and enforce links between data in different tables. They ensure referential integrity by requiring that a value in one table must match a primary key value in a related table. This mechanism prevents inconsistent data entries and maintains logical connections across the database, ensuring that related records remain accurately synchronized and valid.

17. (B) Selection (σ)
The selection operation in relational algebra, denoted by the Greek letter sigma, is used to filter tuples from a relation. It returns only those rows that satisfy a specific condition or predicate. This operation effectively performs horizontal partitioning, allowing users to extract subsets of

A rollback operation is a critical part of transaction management that allows a database to revert to its previous consistent state if an error occurs. When a transaction is rolled back, all changes made since it began are undone, as if the transaction never happened. This ensures the atomicity of transactions, maintaining data integrity by preventing partial or incorrect updates.

55. (A) Both A and R are true and R is the correct explanation of A.
The Two-Phase Locking protocol is a widely used method for ensuring that concurrent transactions are serializable. By requiring that no new locks be acquired once a lock has been released, the protocol prevents the formation of cycles in the precedence graph. The reason accurately explains how this restriction on lock management guarantees that the resulting execution of transactions remains consistent.

56. (B) Normalization to BCNF
Object-oriented databases focus on representing data as complex objects with encapsulation, class hierarchies, and inheritance. While they provide powerful tools for modeling intricate relationships, traditional normalization rules like BCNF are concepts primarily associated with the relational model. Object-oriented systems handle data redundancy and organization through different mechanisms, such as object identity and inheritance, rather than formal relational normalization techniques.

57. (A) Data Control Language (DCL)
Data Control Language is a subset of SQL used primarily to manage access and security within a database. The grant and revoke commands are the core components of DCL, allowing administrators to provide or remove specific privileges for users on various database objects. This ensures that only authorized individuals can perform certain operations, protecting the confidentiality and integrity of stored data.

58. (C) A transaction reads data that has been modified by another transaction that has not yet committed.
A dirty read occurs in a concurrent environment when a transaction accesses data that has been updated by another transaction that has not yet been finalized. If the second transaction is later rolled back, the first transaction will have used invalid, non-existent data. This concurrency problem can lead to significant inconsistencies and is typically managed through different isolation levels in the database.

59. (C) Third Normal Form (3NF)
In the relation R(A, B, C, D, E) with primary key AB, the dependency C to D represents a situation where a non-prime attribute determines another non-prime attribute. This is a transitive dependency, which is permitted in 2NF but strictly forbidden in 3NF. To achieve 3NF, the relation must be decomposed to ensure that all non-prime attributes are dependent only on the primary key.

60. (A) a, b, and c only
Database recovery techniques are essential for restoring a system to a consistent state after a failure. Log-based recovery, shadow paging, and checkpointing are all valid methods used for this purpose. While two-phase locking and timestamp ordering are important database concepts, they are techniques used for concurrency control to manage simultaneous transactions rather than for recovering data after a system crash.

61. (D) A database that stores data relating to time instances, associating a specific time period with each stored tuple.
A temporal database is specifically designed to manage data that changes over time by associating time periods with stored information. This allows the system to keep track of historical changes and query the state of data at any point in the past. Unlike traditional databases that only store the current state, temporal systems provide a comprehensive history of all data modifications.

62. (B) Deductive Database
A deductive database combines the capabilities of a traditional database with the power of logic programming. It uses a set of facts along with a collection of rules to infer new information that is not explicitly stored. This ability to deduce new facts from existing data makes it particularly useful for applications involving complex logical relationships and sophisticated automated reasoning tasks.

Click Here to Buy book on Amazon

the new column name and its corresponding data type.

27. (D) Variable-length character strings.
The VARCHAR data type in SQL is specifically designed to store character strings of varying lengths. Unlike fixed-length character types, VARCHAR only uses as much storage space as needed for the actual text provided, plus a small amount of overhead. This makes it an efficient choice for storing data like names or descriptions where the length of entries can vary significantly.

28. (B) To provide a virtual table based on the result-set of an SQL statement.
An SQL view is a virtual table that does not store data physically but presents it based on a saved query. It allows users to simplify complex queries, restrict access to specific columns, and present data from multiple tables as a single entity. Views provide a layer of abstraction that enhances security and simplifies the user interaction with the database.

29. (C) Unauthorized access to sensitive data.
SQL injection is a severe security vulnerability that occurs when malicious SQL code is inserted into input fields. This can trick the database into executing unintended commands, leading to unauthorized access, data theft, or even the deletion of entire tables. Protecting against these attacks is critical for maintaining the confidentiality and integrity of sensitive information stored within a database system.

30. (A) a, b, and d only
Triggers are specialized stored procedures that execute automatically in response to database events like insert, update, or delete. They are powerful tools for maintaining data integrity and enforcing business rules. While they can be defined to run before or after an operation, they are also capable of calling other stored procedures, making the third statement in the provided list incorrect.

31. (A) `HAVING`
In SQL queries, the having clause is used to filter results after an aggregation has been performed. While the where clause filters individual rows before grouping, having is applied to the groups themselves based on aggregate functions like sum or count. This allows users to retrieve only those grouped results that satisfy specific conditions involving summary data across multiple database records.

32. (A) Both Statement I and Statement II are correct.
Stored procedures are reusable blocks of SQL code saved in the database for efficiency and consistency. They can perform complex operations and do not require a return value. Conversely, SQL functions are designed to return either a single value or a table, often for use within other queries. Both statements correctly describe the fundamental differences and purposes of these database objects.

33. (B) `DROP VIEW view_name`
To permanently remove a view from a database, the drop view command is used followed by the name of the specific view. This is a Data Definition Language command that deletes the view's definition from the database schema. Unlike deleting data from a table, dropping a view only removes the virtual table definition and does not affect the underlying data.

NODIA

34. (A) Using prepared statements with parameterized queries
Prepared statements with parameterized queries are the most effective way to prevent SQL injection. This technique separates the SQL command from the user-provided data, ensuring that input is always treated as a literal value rather than executable code. By pre-compiling the SQL structure, the database can safely handle user input, significantly reducing the risk of malicious command execution.

35. (A) Both A and R are true and R is the correct explanation of A.
Normalization is a critical process in relational database design that involves organizing data to minimize redundancy. By decomposing tables into smaller, well-structured units, normalization helps prevent update, insertion, and deletion anomalies. The provided reason correctly explains how this systematic approach leads to improved data integrity and a more efficient database structure, validating the importance of the normalization process in design.

36. (D) `CREATE`

SQL commands are categorized based on their primary function within the database system. Data Manipulation Language commands like insert, update, and delete are used to modify the actual data stored in tables. In contrast, create is a Data Definition Language command used to define or alter the database structure itself. Therefore, create is not classified as a manipulation command.

37. (C) It ensures that all values in a column are different, but allows one NULL value.
The unique constraint in SQL is used to ensure that all values in a specific column are distinct from each other. Unlike the primary key constraint, which also requires uniqueness but forbids null values entirely, a unique constraint typically allows a single null value to exist in the column. This constraint is essential for maintaining data integrity across non-primary identification fields.

38. (B) Trigger
A trigger is a specialized database object that the management system automatically executes when a specific event occurs, such as an insert, update, or delete. These are often used to enforce complex business rules, maintain audit trails, or synchronize data between related tables. Triggers operate behind the scenes, ensuring that the defined actions are taken consistently whenever the specified events occur.

39. (D) A trigger on the `Orders` table
In an online store database, a trigger on the orders table is the most effective way to automate inventory updates. When a new record is inserted into the orders table, the trigger can be programmed to automatically decrease the corresponding product count in the inventory. This ensures that stock levels are updated in real-time without requiring manual intervention from application code.

40. (C) a, b, c, and d only
SQL is a declarative language where users specify what data to retrieve rather than how to get it. Commands are categorized into DDL for structure, DML for data manipulation, and DCL for access control. While keywords are generally case-insensitive, these categories provide a structured framework for database interaction. Statement e is incorrect because SQL keywords are not case-sensitive.

41. (D) For any two tuples t1 and t2 in R that have t1[X] = t2[X], they must also have t1[Y] = t2[Y].
A functional dependency exists between two sets of attributes when the value of one set uniquely determines the value of another. In a relation, if two rows have the same values for the attribute set X, they must also have identical values for the attribute set Y. This principle is fundamental for understanding database normalization and ensuring data consistency within tables.

42. (B) The domain of each attribute contains only atomic values.
A relation is considered to be in First Normal Form if every attribute contains only atomic, or indivisible, values. This means that multi-valued attributes and nested tables are not permitted within the structure. By ensuring that each field contains only a single value from the attribute's domain, 1NF provides a consistent and simplified structure for data storage and retrieval in databases.

43. (C) Durability: The changes applied to the database by a committed transaction can be lost in case of a system failure.
The durability property of a transaction guarantees that once a transaction has been committed, its changes are permanently recorded and will survive any subsequent system failures. The provided option is incorrect because it suggests that committed data can be lost, which directly contradicts the definition of durability. This property ensures that the results of successful operations are always persistent and reliable.

44. (D) Optimistic Concurrency Control (Validation Based)
Optimistic concurrency control is a technique where transactions are allowed to proceed without locking data. Instead, the system checks for conflicts only when a transaction is ready to commit, during a validation phase. If a conflict is detected, the transaction is rolled back. This approach is most effective in environments where data contention is low and transaction interference is relatively rare.

Downloaded from www.pdf.tube

45. (C) P-4, Q-1, R-2, S-3
Different normal forms address various types of redundancy and anomalies in database design. 1NF requires atomic values, while 2NF eliminates partial dependencies on the primary key. 3NF removes transitive dependencies where non-key attributes depend on other non-key attributes. BCNF is a stricter form where every determinant must be a superkey. Together, these rules ensure a highly structured and efficient database.

46. (B) To find the query execution plan with the minimum resource usage.
Query optimization is a core component of a database management system that aims to identify the most efficient way to execute a given SQL statement. The optimizer evaluates various execution plans and selects the one expected to use the fewest resources, such as CPU time and disk I/O. This process is essential for ensuring high performance and fast response times.

47. (D) No non-prime attribute is partially dependent on any candidate key.
A relation is in Second Normal Form if it is already in 1NF and every non-prime attribute is fully functionally dependent on the entire primary key. This means that no non-prime attribute should depend on only a part of a composite candidate key. Eliminating partial dependencies helps to reduce redundancy and prevents potential update anomalies within the database table structure.

48. (A) Transitive dependency
In the given relation R(A, B, C, D) with dependencies A to B and B to C, a transitive dependency exists because A determines C through B. Since A is the primary key, the non-prime attribute C depends on another non-prime attribute B rather than directly on the primary key. This type of dependency is what 3NF aims to eliminate during normalization.

49. (C) Checkpointing with Log-Based Recovery
Database recovery often involves using a log that records all transactions and their changes. When a failure occurs, the system can use this log to redo committed transactions and undo uncommitted ones. Checkpointing is used to periodically synchronize the log with the actual database, reducing the amount of log data that needs to be processed during the recovery operation after a crash.

50. (A) a, b, and d only
Boyce-Codd Normal Form is a stricter version of Third Normal Form that requires every determinant in a non-trivial functional dependency to be a superkey. While every relation in BCNF is automatically in 3NF, the reverse is not always true. This distinction is important for identifying and eliminating subtle redundancies that can remain even after a relation has reached 3NF during normalization.

51. (D) Growing Phase and Shrinking Phase
The Two-Phase Locking protocol ensures serializability in concurrent transactions through two distinct stages. During the growing phase, a transaction can acquire new locks but cannot release any. Once the transaction releases its first lock, it enters the shrinking phase, where it can release existing locks but is strictly forbidden from acquiring any new ones. This structure prevents cycles in transaction dependencies.

52. (A) Both Statement I and Statement II are correct.
Discretionary Access Control allows the owners of data objects to manage access permissions based on user identity. In contrast, Mandatory Access Control is a more rigid system where access is determined by security classifications and labels assigned by a central authority. Both statements correctly define these two fundamental approaches to managing database security and controlling how users interact with sensitive information.

53. (C) Relational and Object-Oriented models
Object-Relational Database Management Systems were developed to combine the structural strengths of the relational model with the advanced features of object-oriented programming. They allow for the storage of complex data types, inheritance, and encapsulation within a traditional table-based framework. This hybrid approach enables developers to model real-world objects more naturally while still benefiting from the robustness of relational database technology.

54. (D) It undoes all the changes made by the transaction.

NODIA

71. (D) Disconnected operation and data synchronization.
Mobile databases are designed for users who may have intermittent network connectivity. A key challenge is supporting disconnected operation, where users can continue to interact with a local copy of the data, followed by synchronization with the central server when a connection is restored. Managing these updates and resolving conflicts during the synchronization process is a critical aspect of mobile database design.

72. (A) Both Statement I and Statement II are correct.
An XML database is specifically designed to store and manage data in XML format. Unlike traditional relational databases that require a rigid, predefined schema, XML databases offer more flexibility and can handle semi-structured data without a strict schema definition. Both statements correctly describe these features, which make XML databases particularly suitable for managing diverse and evolving data formats in web applications.

73. (C) Two-tier architecture
In a two-tier client-server architecture, the system is divided into two parts: the client handles the user interface and presentation, while the server is responsible for data processing and storage. This model is commonly used for smaller applications where the client can directly communicate with the database server. It provides a straightforward way to distribute tasks between user-facing components and the back-end system.

74. (D) To infer new facts from existing data using rules.
The defining feature of a deductive database is its ability to use logic-based rules to infer new facts from stored data. By combining a set of explicitly stored facts with a collection of rules, the system can automatically deduce information that is not directly present in the database. This reasoning capability is what distinguishes deductive systems from traditional, non-deductive database management systems.

75. (B) a, b, and c only
Multimedia, temporal, and geographic databases are all considered enhanced data models because they extend the capabilities of the traditional relational model to handle specialized data types and requirements. Hierarchical databases, however, are an older model that predates the relational model and are not considered an enhancement. Therefore, the first three options represent the correct group of advanced, enhanced database technologies.

76. (B) Short, fast, and atomic operations.
Online Transaction Processing systems are designed to handle large numbers of concurrent, relatively simple transactions. These operations are typically short, fast, and must be atomic to ensure data integrity. Examples include processing bank transfers or updating inventory for a single sale. The primary focus of OLTP is on providing high performance and maintaining accurate, current data for daily business operations.

77. (D) Association Rule Mining
Association rule mining is a data mining technique used to identify relationships and patterns between items in large transactional datasets. It focuses on finding rules that predict when the occurrence of one set of items is likely to lead to the occurrence of another set. This is commonly used in market basket analysis to discover which products customers frequently purchase together.

78. (B) OLTP databases are highly normalized, while OLAP databases (data warehouses) are often denormalized.
A key structural difference between OLTP and OLAP systems lies in their degree of normalization. OLTP systems use highly normalized tables to ensure efficient updates and data integrity for daily transactions. In contrast, OLAP systems often use denormalized structures, such as star or snowflake schemas, to optimize performance for complex analytical queries and large-scale data aggregations in data warehouse environments.

79. (C) The percentage of transactions in the database that contain both X and Y.
In association rule mining, the support of a rule measures how frequently the items involved appear together across the entire available dataset. Specifically, it represents the percentage of all transactions that contain both the antecedent and the consequent. High support indicates that a

Click Here to Buy book on Amazon

63. (C) XML (eXtensible Markup Language)
XML is a versatile and widely used language for representing and exchanging data over the internet. Its self-describing nature and platform independence make it an ideal format for sharing information between different database systems and web applications. In internet databases, XML provides the standardized structure necessary for the seamless integration and communication of data across diverse and distributed computing environments.

64. (D) Managing large object sizes and content-based retrieval.
Multimedia databases must handle data like images, audio, and video, which are much larger and more complex than traditional text-based records. A significant challenge in these systems is managing these large objects efficiently and providing effective content-based retrieval methods. Unlike traditional databases that search by keywords or values, multimedia systems often require sophisticated algorithms to search based on visual or auditory features.

65. (B) P-3, Q-2, R-4, S-1
Enhanced data models are tailored for specific application requirements. Geographic Information Systems manage spatial data for location-based queries, while genome data management handles biological sequences. Mobile databases support disconnected operations for portable devices, and temporal databases track historical records over time. Matching these models to their primary applications highlights the diverse ways database technology is adapted to meet complex, real-world data management challenges.

66. (B) Data is logically interrelated but stored in multiple physically separate locations.
A distributed database consists of multiple logically related database fragments that are physically stored across different sites in a network. While the data is physically separated, it appears to the user as a single, integrated database system. This architecture provides benefits such as increased reliability, improved performance through local processing, and better scalability, but it also introduces complexities in managing data.

67. (D) Simpler recovery protocols due to single-site operation.
Client-server architectures offer several advantages, including distributed processing, enhanced security, and potentially lower hardware costs for clients. However, they do not simplify recovery protocols. In fact, managing recovery in a distributed environment where data and processing are split between clients and servers is typically more complex than in a centralized, single-site operation. Therefore, simplified recovery is not a characteristic advantage.

68. (A) To capture, store, manipulate, analyze, and manage spatial or geographical data.
Geographic Information Systems are specialized database systems designed to handle data that is referenced by geographic coordinates. They allow users to store, analyze, and visualize spatial information, making them essential for applications like urban planning, environmental monitoring, and navigation. GIS provides powerful tools for performing spatial queries and analyzing the relationships between different layers of geographical information within the database.

69. (C) Biological sequence and structure data.
Genome data management involves the storage and analysis of complex biological information, such as DNA sequences, protein structures, and genetic maps. This field requires specialized database systems capable of handling massive amounts of sequential data and supporting sophisticated bioinformatics tools for research. The primary focus is on managing the intricate biological data that is fundamental to modern genomics and molecular biology studies.

70. (A) b, a, c, d
The development of database models has followed a chronological path from simpler to more complex structures. Early hierarchical databases were followed by the widely adopted relational model. Later, object-oriented databases emerged to handle complex data types, and more recently, NoSQL databases gained popularity for managing large-scale, unstructured data. This sequence reflects the ongoing evolution of technology to meet changing data management needs.

Downloaded from www.pdf.tube

finding optimal boundaries between data points. Both statements accurately describe these two important models commonly used in data mining and machine learning.

89. (C) Discovering relationships and connections within data.
Link analysis is a data mining technique that focuses on uncovering and analyzing the connections between different entities in a dataset. It is particularly useful for identifying patterns of interaction, such as relationships in a social network or connections between web pages. By focusing on the links between items, this technique provides valuable insights into the structure and influence within complex systems of information.

90. (A) P-4, Q-3, R-2, S-1
Different data mining techniques are used for various types of analysis. Classification assigns items to predefined categories, while clustering groups data based on inherent similarities. Association rule mining discovers interesting if-then patterns in transactions, and regression is used to predict continuous numerical values. Matching these techniques with their descriptions provides a clear overview of the fundamental tasks performed in modern data mining processes.

91. (D) The study of social structures through the use of network and graph theory.
Social Network Analysis is a specialized data mining technique that uses network and graph theory to investigate social structures. It focuses on analyzing the relationships and interactions between actors, such as individuals or organizations, and identifying patterns of influence and connection within the network. This approach provides a powerful way to understand the social dynamics and structural properties of complex online communities.

92. (B) Association Rule
The discovery that customers who buy bread and butter are also likely to buy milk is a classic example of an association rule. This type of pattern identifies items that frequently occur together in transactional data, such as a supermarket's sales records. Association rules are valuable for businesses to optimize product placement, marketing strategies, and recommendation systems based on observed customer behavior.

93. (C) Regression
Regression analysis is the most suitable technique for predicting future stock prices because it is designed to model the relationships between variables and predict continuous numerical values. By analyzing historical performance and other relevant factors, a regression model can estimate a future value based on identified trends. This makes it an essential tool for financial forecasting and other applications involving numerical prediction.

94. (D) To describe the significant dependencies among variables.
Dependency modeling is a data mining task that aims to discover and describe the significant relationships and dependencies between different variables in a dataset. This involves identifying which variables are strongly related and how they influence each other. By modeling these dependencies, analysts can gain a deeper understanding of the underlying structure of the data and the factors that drive specific outcomes.

95. (A) The number of nearest neighbors to consider for classification.
In the K-Nearest Neighbour algorithm, 'K' represents the specific number of nearest neighbors that the model evaluates when making a classification decision. A new data point is assigned to the category that is most frequent among these K closest neighbors in the feature space. The choice of K is a critical parameter that can significantly affect the accuracy and performance of the model.

96. (B) Second Normal Form (2NF)
Second Normal Form is specifically designed to address and eliminate partial functional dependencies within a relation. A partial dependency occurs when a non-prime attribute depends on only a part of a composite candidate key rather than the entire key. To satisfy 2NF, all non-prime attributes must be fully functionally dependent on the primary key, ensuring a more efficient and redundant-free design.

97. (A) Durability

NODIA

particular relationship occurs often enough to be considered significant for analysis and decision-making within the given database.

80. (A) Both A and R are true and R is the correct explanation of A.
Data warehousing is the process of creating and utilizing a specialized repository for analytical data. A data warehouse is defined as a subject-oriented, integrated, and non-volatile collection of data that supports management's decision-making process. The reason correctly explains what a data warehouse is and how it relates to the overall assertion about the data warehousing process in large organizations.

81. (A) Classification
Classification is a fundamental data mining task that involves assigning individual items to one of several predefined categories or classes. This is a supervised learning technique where a model is trained on a labeled dataset to learn how to categorize new, unseen data points. It is widely used for applications like spam detection, credit scoring, and diagnosing medical conditions based on features.

82. (B) To organize data into different levels of abstraction or detail.
In data warehousing, a concept hierarchy is used to organize values for a specific attribute into multiple levels of detail. This allows users to view and analyze data at different levels of abstraction, such as grouping sales data by day, month, or year. Concept hierarchies are essential for performing drill-down and roll-up operations during online analytical processing tasks within the warehouse.

83. (C) Classification
The K-Nearest Neighbour algorithm is a popular supervised learning technique used for classification. It works by identifying the K most similar data points in the training set and assigning a new instance to the class most common among its neighbors. KNN is a simple yet effective method for categorizing data based on the proximity of features in a multidimensional data space.

84. (D) Clustering
Clustering is an unsupervised data mining technique used to group data points based on their similarities without using any predefined labels. This makes it the most appropriate choice for segmenting customers into groups based on their purchasing behavior when those segments are not already known. By identifying natural patterns in the data, clustering helps companies discover meaningful groups within their customer base.

85. (C) a, b, c, and d only
Data warehouses are characterized by being subject-oriented, integrated, time-variant, and non-volatile. These features make them ideal for long-term analytical processing and decision support. However, they are not primarily designed for online transaction processing, which is the domain of traditional relational databases. Therefore, the first four statements accurately describe the key properties of a data warehouse, while the fifth statement is incorrect.

86. (B) Normalization
Classification, clustering, and regression are all common data mining techniques used to discover patterns and predict outcomes from large datasets. Normalization, however, is a process used in database design to organize tables and minimize redundancy. While normalization is a critical part of database management, it is not considered a data mining technique for extracting knowledge or patterns from information.

87. (C) To predict a continuous-valued output for a given input.
Regression is a statistical data mining technique used to model the relationship between a dependent variable and one or more independent variables. Its primary goal is to predict a continuous numerical value, such as a company's future sales or the price of a house, based on a set of input features. This distinguishes it from classification, which predicts categorical labels.

88. (A) Both Statement I and Statement II are correct.
A Hidden Markov Model is a statistical tool used for modeling sequences of events where the underlying states are not directly observable. Support Vector Machines are a powerful supervised learning model used for both classification and regression tasks by

Click Here to Buy book on Amazon

data based on defined criteria without altering the structure or columns of the original relation being queried.

18. (A) Both Statement I and Statement II are correct.
Tuple Relational Calculus focuses on identifying entire tuples that satisfy a given predicate. In contrast, Domain Relational Calculus uses variables that range over the individual domains of attributes rather than entire rows. Both statements are correct because they accurately describe the two different approaches used in formal non-procedural query languages to specify what data should be retrieved from a database.

19. (C) Graphical User Interface Rule
Codd's twelve rules define the requirements for a management system to be considered fully relational. While the Information Rule, View Updating Rule, and Logical Data Independence are core components of this framework, there is no Graphical User Interface Rule. These rules focus on data integrity, manipulation, and structural independence rather than the specific visual design of user interface components.

20. (A) Both A and R are true and R is the correct explanation of A.
Relational databases rely on integrity constraints like primary keys, foreign keys, and domain constraints to ensure data accuracy and consistency. Consequently, all update operations must be carefully managed to handle potential violations of these rules. The provided reason correctly explains why the assertion is true, as enforcing these constraints is necessary to prevent the entry of invalid or inconsistent data.

21. (C) Reject the operation.
When an operation like an insert violates an integrity constraint, the most common default action for a database management system is to reject the operation entirely. This prevents the database from entering an inconsistent state. By refusing to accept the invalid data, the system ensures that all stored information strictly adheres to the predefined rules and relationships within the database schema.

22. (B) Projection
The projection operation in relational algebra is used to select specific attributes or columns from a relation. It allows users to retrieve only the data fields they are interested in, while discarding others. This operation performs a vertical partitioning of the table, effectively reducing the width of the relation to include only the specified set of attributes in the output results.

23. (D) TRC variables range over tuples, while DRC variables range over domains of attributes.
Tuple Relational Calculus and Domain Relational Calculus are both non-procedural formal languages used in database theory. The primary difference between them lies in the range of their variables. In TRC, variables represent entire tuples within a relation, whereas in DRC, variables represent individual values from the domains of attributes. Both languages provide equivalent expressive power for querying relational database systems.

24. (A) One-to-Many (1:N)
In a university database design, the relationship from professor to student is classified as one-to-many. This cardinality indicates that a single professor can act as an advisor to multiple students. However, the requirement specifies that each student can have only one advisor. Therefore, the mapping from the professor entity set to the student entity set follows the one-to-many structural pattern.

25. (B) P-4, Q-1, R-2, S-3
Entity integrity ensures primary keys are not null, while referential integrity requires foreign keys to match existing primary keys or be null. Domain constraints specify valid attribute values, and key constraints prevent duplicate superkey values. Matching these definitions accurately describes how relational databases use different types of constraints to maintain the overall quality, consistency, and structural integrity of the stored information.

26. (A) `ALTER TABLE table_name ADD COLUMN new_col_name datatype;`
The SQL language provides specific commands to modify the structure of existing tables without losing the data already stored within them. The correct syntax for adding a new field is to use the alter table command followed by the add column clause. This allows developers to extend the database schema dynamically by specifying

NODIA

The durability property of a database transaction ensures that once it has been successfully committed, its changes are permanent and will not be lost, even if a system failure occurs. This guarantee is essential for maintaining the reliability and persistence of the data. Durability is typically achieved through logging and other recovery mechanisms that allow the system to reconstruct committed changes after a crash.

98. (B) Database recovery
A log record is a critical component of a database system's recovery mechanism. It contains detailed information about all modifications made to the database, including the start and end of transactions and the specific changes to data items. In the event of a system failure, the recovery manager uses these log records to restore the database to a consistent state by redoing or undoing operations.

99. (C) The pure relational model with no extensions
Object-relational databases were created to extend the traditional relational model with advanced features like user-defined data types, inheritance, and support for complex objects. By definition, an ORDBMS goes beyond the pure relational model by incorporating these object-oriented capabilities. Therefore, the pure relational model without any extensions is not what an ORDBMS aims to support; instead, it provides a more versatile hybrid framework.

100. (D) a, b, c, and d
Query processing and optimization are essential for the efficient operation of a database management system. Optimization aims to minimize execution time by using cost estimates and relational algebra equivalences to find the most efficient execution plan. All the provided statements correctly describe different aspects of this process, highlighting the importance of statistics and logical transformations in achieving high performance for database queries.

101. (D) Shadow paging
Timestamp ordering, two-phase locking, and validation techniques are all recognized methods used for concurrency control to manage simultaneous transactions. Shadow paging, however, is a technique used for database recovery and ensuring transaction atomicity rather than for managing concurrent access. While both are important parts of a DBMS, shadow paging serves a different primary purpose related to recovering from failures.

102. (A) Both Statement I and Statement II are correct.
Normalization is a valuable process for eliminating data anomalies and ensuring integrity in a relational database. However, excessive normalization can lead to poorer query performance because it often requires a greater number of joins to retrieve related information. Both statements accurately describe the benefits and potential trade-offs associated with applying different levels of normalization during the database design process.

103. (C) Give privileges to a user.
The grant command is a key part of the Data Control Language in SQL, used by database administrators to provide specific access permissions to users or roles. These privileges can include the ability to select, insert, update, or delete data from specific tables. Granting permissions is a fundamental part of managing database security and ensuring that users only have access to the information they need.

104. (A) Phantom Read
A phantom read is a concurrency problem that occurs when a transaction retrieves a set of rows satisfying a condition, but a subsequent query within the same transaction returns a different set of rows because another transaction has concurrently inserted or deleted data. This phenomenon can lead to inconsistencies in reports and analysis and is typically managed through the highest database isolation levels.

105. (C) a, b, c, and d only
Transactions, concurrency control, database recovery, and functional dependency are all fundamental concepts in database management. The first four statements accurately define these terms and their roles within a system. However, the fifth statement is incorrect because BCNF is actually a stricter normal form than 3NF. Therefore, the first four options represent the set of correct and factual statements about database systems.

Downloaded from www.pdf.tube

176. (B) Clustering
Clustering is a data mining task that is most similar to unsupervised learning because it involves finding natural groups in data without using any predefined labels or classes. Like other unsupervised methods, clustering focuses on identifying inherent patterns and similarities within the dataset itself. This distinguishes it from classification, which is a supervised learning task that relies on labeled training data to categorize new, unseen information.

177. (C) Data Cube
The conceptual model for Online Analytical Processing is often represented as a data cube, which allows for the multidimensional analysis of information. A data cube organizes data along several dimensions, such as time, product, and region, making it easy to perform complex aggregations and drill-down operations. This model provides an intuitive way for users to visualize and explore large datasets from many different analytical perspectives simultaneously.

178. (D) It measures the reliability of the rule, indicating how often the given rule is found to be accurate.
In association rule mining, confidence measures the reliability or strength of a rule by indicating how often the consequent appears in transactions that also contain the antecedent. It is calculated as the proportion of transactions containing both items compared to the total number of transactions containing the antecedent. High confidence suggests a strong and reliable relationship between the items, making the rule more valuable for prediction.

179. (C) Does not conform to a traditional row-column database structure.
Unstructured data is a hallmark of Big Data and is defined by its lack of a predefined data model or organizational structure. Unlike structured data, it does not fit easily into the rows and columns of a traditional relational database. Examples include text documents, images, videos, and social media posts. Managing and analyzing this diverse and unstructured information requires specialized tools and architectures beyond traditional database management systems.

180. (B) P-4, Q-3, R-1, S-2
Different SQL objects serve distinct purposes in a database system. A trigger is automatically executed in response to an event, while a view provides a virtual table based on a query. A stored procedure is a reusable block of code, and a constraint is a rule used to maintain data integrity. Matching these terms to their purposes provides a clear overview of the various tools available for database development and management.

181. (B) Cassandra
Cassandra is a highly scalable and distributed NoSQL database modeled after Google's Bigtable. It is organized as a multi-dimensional sorted map and is designed to handle massive amounts of data across many commodity servers. Cassandra provides high availability with no single point of failure, making it an excellent choice for large-scale applications that require continuous operation and fast performance in a distributed environment.

182. (A) Both Statement I and Statement II are correct.
Data independence is the vital ability to modify a schema at one level without affecting the next higher level in a database system. Additionally, the client-server architecture is a specific form of distributed system where tasks are split between different machines. Both statements accurately describe these fundamental database concepts, which are essential for creating flexible, scalable, and maintainable systems for managing large collections of data.

183. (D) Ellipses
In an Entity-Relationship diagram, ellipses are the standard graphical notation used to represent attributes. These ellipses are linked to the entities or relationships they describe, providing a clear visual way to show the properties and characteristics of each component in the model. This consistent visual language helps database designers communicate and document the structure of the data clearly and effectively during the design process.

184. (B) SET NULL / SET DEFAULT
When a foreign key constraint is violated during a deletion or update, setting the foreign key value to null or a predefined default value is one possible remedial action. These options allow the database to maintain referential integrity

NODIA

organize and prepare massive amounts of data for final aggregation.

115. (C) Block
In the Hadoop Distributed File System, files are broken down into smaller, fixed-size units called blocks. These blocks are the fundamental units of data storage and are replicated across different nodes in the cluster to ensure fault tolerance and high availability. By managing data at the block level, HDFS can efficiently handle extremely large files and provide reliable storage in a distributed environment.

116. (D) They provide storage and retrieval for data modeled by means other than tabular relations used in RDBMS.
NoSQL databases are a category of data management systems designed to handle data that does not fit easily into the traditional tabular structure of relational databases. They use various data models, such as document, key-value, and graph, to provide more flexibility and scalability. NoSQL systems are particularly useful for large-scale applications with rapidly changing data requirements and high throughput needs in distributed environments.

117. (B) MongoDB
MongoDB is a popular open-source NoSQL database that uses a document-oriented data model. It stores information in flexible, JSON-like documents, allowing for a dynamic and evolving schema. This distinguishes it from traditional relational databases like MySQL, Oracle, and SQL Server, which use a rigid tabular structure. MongoDB is widely used for modern web applications that require high scalability and flexibility.

118. (A) They may support SQL-like query languages but are not based on the relational model.
The term NoSQL is often interpreted as Not Only SQL, implying that while these databases are not based on the traditional relational model, they may still support SQL-like query languages. This means that they provide alternative ways to store and retrieve data while sometimes offering familiar querying syntax. This flexibility allows developers to choose the best data model for their specific application needs.

119. (C) Horizontal scalability and flexibility.
NoSQL databases are particularly well-suited for cloud environments because of their ability to scale horizontally by adding more commodity servers to a cluster. This allows them to handle increasing amounts of data and traffic with ease. Additionally, their flexible schema design makes it simple to adapt to changing data requirements, which is essential for the dynamic and rapidly evolving nature of modern applications.

120. (B) P-3, Q-1, R-4, S-2
NoSQL databases come in several different types, each with its own unique data model and characteristics. Key-value stores are the simplest, using a hash table approach. Document stores use JSON or BSON formats, while column-family stores organize data into dynamic columns. Graph databases are specialized for traversing complex relationships. Matching these types to their core characteristics provides a clear overview of the diverse NoSQL landscape.

121. (D) Minimizing the number of joins, as they are often expensive or unsupported.
In NoSQL systems, query optimization often focuses on minimizing or avoiding join operations, which are typically expensive in distributed environments and sometimes not supported at all. Instead, data is often denormalized or structured in a way that allows for efficient retrieval without needing to combine information from multiple collections. This approach prioritizes read performance and horizontal scalability over traditional relational normalization rules.

122. (A) Both Statement I and Statement II are correct.
NoSQL databases often prioritize scalability and availability over strict consistency, a trade-off described by the BASE properties. Additionally, while indexing is critical for performance in both NoSQL and relational systems, the implementation details vary significantly based on the specific NoSQL data model being used. Both statements accurately reflect these key considerations in the design and operation of modern NoSQL database management systems.

123. (C) Graph Database
A graph database is the most suitable choice for social media applications where connections and

Click Here to Buy book on Amazon

A temporal database is specifically designed to handle data that changes over time, allowing for queries about historical states. A system tracking employee salaries over two decades and supporting queries about specific past years is a perfect example of a temporal application. These databases provide the specialized structures and operators necessary for managing and querying the evolution of data through time.

142. (B) eXtensible Markup Language
XML, or eXtensible Markup Language, is a standardized format for representing and sharing structured data over the internet. Its name reflects its primary characteristic: the ability for users to define their own tags, making it highly flexible and extensible for diverse data requirements. In the context of internet databases, XML provides a common language for exchanging information between different systems and platforms.

143. (C) Single point of failure
Distributed databases are designed to provide transparency in terms of location, replication, and fragmentation, allowing users to interact with data without knowing its physical distribution. A key goal is to improve reliability by avoiding a single point of failure. Therefore, having such a point of failure is not a feature of a well-designed distributed database system; instead, they aim for resilience through data distribution.

144. (D) Regression
Regression is the most appropriate data mining technique for predicting a continuous numerical value, such as a student's final grade. By analyzing the relationship between the final grade and independent variables like midterm scores and attendance, a regression model can be used to estimate future outcomes. This makes it an essential tool for forecasting and prediction in many different fields, including education and business.

145. (A) Fact table
In the star schema of a data warehouse, the central table is known as the fact table. It typically contains quantitative data about business processes, such as sales amounts or quantities sold, and contains foreign keys that link it to several dimension tables. The fact table is the core of the star schema, providing the primary data for analysis and aggregation in data warehousing applications.

146. (B) The data is not subject to frequent changes and deletions as in an OLTP system.
A data warehouse is described as non-volatile because once data is loaded, it is typically not updated or deleted in the same way as in a transactional OLTP system. Instead, new data is added as historical snapshots, preserving a consistent record over time. This non-volatile nature is essential for providing a stable and reliable foundation for long-term data analysis and historical reporting.

147. (C) Sequencing Analysis
Sequencing analysis is a data mining method used to identify patterns where events follow one another in a specific order over time. This is particularly useful for understanding customer journeys, analyzing scientific data, or predicting future events based on observed sequences. By focusing on the temporal order of events, sequencing analysis provides valuable insights into the dynamic behavior and trends within a dataset.

148. (A) It is a supervised learning algorithm that finds an optimal hyperplane to separate classes.
Support Vector Machines are a powerful supervised learning algorithm used primarily for classification and regression tasks. The core idea is to find an optimal hyperplane that best separates different classes in a high-dimensional feature space. By maximizing the margin between the boundary and the closest data points, SVMs provide a robust and accurate model for categorizing new instances based on learned patterns.

149. (C) Hadoop Distributed File System
HDFS stands for Hadoop Distributed File System, which is the primary storage component of the Hadoop ecosystem. It is designed to store massive amounts of data across multiple nodes in a distributed cluster, providing high throughput and fault tolerance. HDFS's architecture is essential for handling the scale of Big Data, allowing for efficient and reliable storage and processing in large-scale environments.

150. (A) P-3, Q-4, R-1, S-2

NODIA

Codd's rule of guaranteed access ensures that every individual data element in a relational database can be uniquely identified and retrieved without ambiguity. These rules collectively serve as a rigorous benchmark for determining whether a database management system can be considered truly relational. Both statements accurately describe the importance and role of Codd's rules in the field of relational database theory.

133. (C) `SELECT`
The select statement is the core command in SQL for retrieving information from a database. It allows users to specify which columns to return, which tables to query, and which conditions must be met for the data to be included in the result set. As the most frequently used SQL keyword, select is fundamental for data analysis and reporting in relational systems.

134. (D) Stored Procedure
A stored procedure is a named, precompiled block of SQL statements that can accept input parameters and perform a series of operations within the database. Unlike functions, they are not required to return a value. Stored procedures are useful for encapsulating complex logic, improving performance through precompilation, and ensuring that specific tasks are executed consistently across different applications and user sessions.

135. (C) a, b, c, and d only
Protecting against SQL injection involves several important security practices. Using parameterized queries, validating input, employing stored procedures, and applying the principle of least privilege are all effective and recommended remedies. Simply disabling all input forms is not a practical solution for most applications. Therefore, the first four statements represent the correct set of strategies for building secure and robust database-driven systems.

136. (B) A determines B.
A functional dependency A to B means that for every unique value of A, there is exactly one corresponding value of B. In this relationship, A is said to determine B. This concept is fundamental for database normalization, as it helps designers identify how attributes are logically related and how to structure tables to avoid redundancy and maintain data consistency within the database.

137. (A) All of the above
A query optimizer evaluates several different cost components to find the most efficient execution plan for an SQL statement. These include the CPU time required for processing, the number of I/O operations for reading and writing data, and the network communication cost for distributed queries. By considering all of these factors, the optimizer can choose the plan that minimizes overall resource usage.

138. (B) Non-repeatable Read
A non-repeatable read is a concurrency problem that occurs when a transaction reads a data item and then reads it again later, only to find that the value has changed because another transaction has committed an update in the meantime. This inconsistency can cause errors in long-running transactions and is managed by using isolation levels that provide more consistent views of the data.

139. (C) To reduce the amount of log data that needs to be scanned during recovery.
A checkpoint is a periodic event where the database management system synchronizes its internal buffers with the data on disk and records the event in the log. Its primary purpose is to limit the amount of work the system must perform during recovery after a failure. By knowing that all changes prior to the checkpoint are safely on disk, the recovery manager can scan only the subsequent log records.

NODIA

140. (A) Both A and R are true and R is the correct explanation of A.
Mandatory Access Control is generally considered more secure than Discretionary Access Control because it enforces policies that cannot be changed by individual users. MAC policies are determined by a central authority and based on security labels, which prevents users from accidentally or intentionally granting unauthorized access to sensitive data. The reason correctly explains how this centralized control provides a higher level of security.

141. (D) Temporal Database

relationships between users are complex and frequently queried. By modeling users as nodes and connections as edges, graph databases allow for extremely efficient traversal of relationships, such as finding friends of friends. This specialized structure provides much better performance for these types of queries than traditional or other NoSQL models.

124. (D) Using natural sort order of primary keys or clustering keys
To improve query performance, NoSQL databases often rely on the natural sort order of primary keys or specialized clustering keys. By storing data in a specific order on disk, the system can more efficiently perform range scans and retrieve related records. This strategy is essential for optimizing read operations in large-scale, distributed environments where traditional indexing and normalization techniques may be less effective.

125. (A) b, a, d, c
The widespread adoption of different data storage paradigms has followed a logical progression over time. Early file systems were followed by the dominance of relational databases. Object-oriented databases then emerged to handle complex data, and more recently, NoSQL databases gained significant traction for Big Data applications. This chronological sequence reflects how technology has evolved to meet increasingly sophisticated and large-scale data management requirements.

126. (B) Data Definition Language (DDL)
Creating a new table in a database is a structural change, which falls under the category of Data Definition Language. DDL commands are specifically used by database administrators and developers to define, modify, and delete the objects that make up the database schema, such as tables, indexes, and views. This language provides the essential commands for establishing the logical framework of the data repository.

127. (C) Data Independence : Tightly couples the schema levels.
Data independence is a key goal of the three-schema architecture, but its purpose is to decouple the schema levels rather than tightly couple them. By allowing changes at one level without affecting others, data independence provides flexibility and simplifies maintenance. All the other pairs correctly describe the different levels and goals of this architecture, making the third option the only incorrect match.

128. (B) The physical distribution of DBMS components.
The main difference between centralized and client-server DBMS architectures is how the physical components of the system are distributed. In a centralized model, everything resides on a single computer. In a client-server model, the database functionality is located on a server, while the user interface and application logic are on client machines. This physical distribution of tasks is what defines the two different architectural paradigms.

129. (A) π (Projection)
In relational algebra, the projection operation, denoted by the Greek letter pi, is used to extract specific columns from a relation. It allows users to focus on the particular attributes they need, effectively reducing the width of the table in the final output. This operation is fundamental for data retrieval, enabling users to customize the set of information returned by a query.

130. (C) a, c, and d only
When a constraint violation occurs during an update, a DBMS can take several valid actions. It can reject the operation to maintain integrity, trigger a user-defined error-correction routine, or cascade the update to related tables in the case of referential integrity. Simply performing the update and logging the violation is generally not a valid action, as it would leave the database inconsistent.

131. (D) A double ellipse
In an Entity-Relationship diagram, attributes are typically represented by ellipses. However, if an attribute can have multiple values for a single entity instance, such as multiple phone numbers for an employee, it is called a multi-valued attribute and is specifically represented by a double ellipse. this visual notation helps designers clearly identify attributes that need to be handled carefully during database design.

132. (A) Both Statement I and Statement II are correct.

Downloaded from www.pdf.tube

Big Data is characterized by four key dimensions: volume, velocity, variety, and veracity. Volume refers to the scale of data, velocity to the speed of processing, variety to the different types of data, and veracity to the trustworthiness of the information. Matching these characteristics to their descriptions provides a comprehensive overview of the fundamental challenges and properties associated with managing and analyzing Big Data.

151. (C) Aggregating the intermediate key-value pairs to produce the final output.
In a MapReduce job, the reduce phase is responsible for aggregating and summarizing the intermediate key-value pairs generated by the previous map phase. It processes each group of values associated with a specific key to produce the final output for the task. This step is essential for distilling massive amounts of processed data into a concise and useful result for the user.

152. (B) Document Store
A document store NoSQL database, such as MongoDB, is the best choice for a content management system with articles of varying attributes. Its flexible, schema-less design allows each article document to have its own unique set of fields, such as author, title, and multiple tags. This flexibility makes it much easier to manage diverse and evolving content than the rigid structure of a relational database.

153. (C) Greater horizontal scalability and handling of unstructured data.
The primary motivation for developing NoSQL databases was the need for systems that could scale horizontally and handle large volumes of unstructured or semi-structured data. Traditional relational databases often struggle with these requirements, especially in large-scale web and Big Data environments. NoSQL systems provide more flexible data models and architectures that are better suited for the diverse and rapidly changing data needs of modern applications.

154. (D) Improve the speed of data retrieval operations.
Indexing in NoSQL databases, just like in relational systems, is primarily used to improve the performance and speed of data retrieval. By creating indexes on frequently queried fields, the system can more quickly locate and access the relevant records, significantly reducing the time required for queries. This is essential for maintaining high performance in large-scale applications with massive amounts of data stored in NoSQL collections.

155. (A) External Schema or View
An external schema, also known as a view, is a logical description of a specific portion of the database that is tailored for a particular group of users. It provides a customized perspective on the data, hiding unnecessary details and providing a simplified interface for interaction. This abstraction is a key part of the three-schema architecture, helping to ensure security and simplify the user experience.

156. (B) Logical Data Independence
Logical data independence is the ability to modify the conceptual schema without requiring changes to the external schemas or the application programs that use them. This allows database administrators to restructure the logical design of the database, such as adding or removing attributes, without impacting the views provided to users. It is an important feature that provides flexibility and simplifies the maintenance of complex systems.

157. (C) Transitive Closure
Union, difference, and Cartesian product are all fundamental operations in relational algebra used to manipulate and combine relations. Transitive closure, however, is not a standard relational algebra operator. While it is an important concept in graph theory and some advanced database languages, it is not part of the basic set of operators used in the formal relational algebra framework for querying relational databases.

158. (D) When a referenced row is deleted, all corresponding referencing rows are also deleted.
The on delete cascade action in a foreign key constraint ensures that if a row in the parent table is deleted, all corresponding rows in the child table are also automatically removed. This is a powerful tool for maintaining referential integrity and ensuring that the database does not contain orphaned records. However, it should be used carefully to avoid unintended data loss when deleting records.

NODIA

106. (B) R has a non-trivial functional dependency X → A where X is not a superkey, and A is a prime attribute.
A relation that is in 3NF but not in BCNF must contain a non-trivial functional dependency where the determinant is not a superkey. The key difference is that 3NF permits such a dependency if the dependent attribute is a prime attribute, while BCNF forbids it entirely. This subtle distinction allows for certain redundancies in 3NF that are eliminated in the stricter BCNF.

107. (D) Query Optimization : Functional Dependency
Checkpointing is a key part of recovery, two-phase locking is a concurrency control method, and grant and revoke are core security functions. These are all correctly matched. However, functional dependency is a concept used in database normalization to organize tables, not a technique for query optimization. Therefore, this final pair is incorrectly matched as it links two unrelated database design and performance concepts.

108. (A) Volume, Velocity, and Variety
Big Data is commonly characterized by three Vs: volume, velocity, and variety. Volume refers to the massive amount of data being generated, velocity to the high speed at which it is created and processed, and variety to the many different formats and sources of information. Together, these factors represent the key challenges that distinguish Big Data from traditional database management environments.

109. (C) A framework for processing large data sets in parallel across a cluster.
MapReduce is a core component of the Hadoop ecosystem, providing a framework for processing massive amounts of data in parallel across a distributed cluster of computers. It works by breaking down a large task into smaller sub-tasks that are executed simultaneously, significantly reducing the time required for processing. This parallel approach is essential for handling the scale and complexity of Big Data.

110. (A) Both A and R are true and R is the correct explanation of A.
The Hadoop Distributed File System is designed for high fault tolerance, ensuring that data remains available even if individual nodes in the cluster fail. It achieves this by automatically replicating data blocks across multiple different machines. The reason correctly explains how this replication mechanism provides the fault tolerance described in the assertion, making HDFS a reliable storage solution for large-scale Big Data applications.

111. (D) Normalized Data
Big Data can exist in three primary forms: structured, semi-structured, and unstructured. Structured data is highly organized, while unstructured data has no predefined model. Semi-structured data falls somewhere in between. Normalized data is a concept from relational database design used to organize information and reduce redundancy, but it is not a category of data types in the context of Big Data classification.

112. (B) It is an open-source framework for storing and processing Big Data.
Hadoop is a widely used open-source framework specifically designed for the distributed storage and processing of massive Big Data sets. It allows organizations to manage data that is too large or complex for traditional relational databases by using clusters of commodity hardware. Hadoop's architecture provides a scalable and cost-effective solution for handling the volume, velocity, and variety of modern data.

113. (A) The speed at which data is generated and must be processed.
The velocity characteristic of Big Data refers to the high speed at which data is continuously generated and the need for it to be processed in a timely manner. This includes real-time data streams from social media, sensors, and transactional systems. Managing this rapid flow of information is a significant challenge, requiring specialized tools and architectures for fast ingestion and analysis.

114. (B) Filtering and sorting
In a MapReduce job, the map phase is the first step where input data is processed to generate intermediate key-value pairs. This typically involves filtering the data based on certain criteria and sorting it to prepare for the subsequent reduce phase. By performing these operations in parallel across many nodes, MapReduce can efficiently

Click Here to Buy book on Amazon

and are widely used in applications like speech recognition and bioinformatics. By modeling the transitions between hidden states and the probability of observations, HMMs provide a powerful framework for understanding complex temporal patterns in various datasets.

168. (A) To produce a compact and concise description of a dataset.
The primary goal of summarization in data mining is to provide a concise and understandable overview of a large and complex dataset. This can involve calculating descriptive statistics, generating summary tables, or identifying the key characteristics of the data. Summarization helps analysts quickly grasp the main features and trends within the information, facilitating easier interpretation and more effective decision-making based on the data.

169. (C) It manages the file system namespace and metadata for all files and directories.
In the Hadoop Distributed File System, the NameNode is a critical component that acts as the master for the cluster. It is responsible for managing the file system namespace, including the structure of directories and the metadata for all files and blocks. Without the NameNode, the system would not know where the actual data blocks are stored, making it an essential part of the distributed storage architecture.

170. (A) Both A and R are true and R is the correct explanation of A.
NoSQL databases are often a good choice for handling large-scale, unstructured data and high-throughput requirements. This is because they frequently prioritize performance, scalability, and availability over strict ACID consistency, following the BASE model instead. The reason accurately explains this common trade-off, clarifying why NoSQL systems are well-suited for modern web applications that need to manage massive amounts of diverse and rapidly changing data.

171. (C) Non-procedural or Declarative
A non-procedural or declarative query language allows users to specify what data they want to retrieve without describing the step-by-step procedure for how to get it. SQL is the most common example, where users write queries that define the desired result set, and the database management system determines the most efficient way to execute them. This abstraction simplifies database interaction and allows for better optimization by the system.

172. (D) Superkey
A relation is considered to be in Boyce-Codd Normal Form if and only if for every non-trivial functional dependency, the determinant attribute set is a superkey. This is a stricter requirement than 3NF and ensures that no attribute, prime or non-prime, is determined by a set of attributes that is not a superkey. Reaching BCNF helps to eliminate subtle redundancies and potential anomalies in relational database design.

NODIA

173. (A) Atomicity
The A in the ACID properties of a database transaction stands for atomicity. This property guarantees that a transaction is treated as a single, indivisible unit of work, which is either completed entirely or not at all. If any part of the transaction fails, the entire operation is rolled back, ensuring that the database never remains in an inconsistent, partially updated state due to an error.

174. (B) Spatial data representation and querying
A core component of any Geographic Information System is its ability to represent and query spatial or geographical data. GIS provides specialized data types and operators for handling coordinates, shapes, and the spatial relationships between them. This capability is what distinguishes GIS from traditional database systems and makes it an essential tool for applications involving maps, urban planning, environmental analysis, and other location-based services.

175. (A) c, d, b, a
Database design follows a structured logical sequence. It begins with conceptual design to create a high-level model, followed by logical design to map this to a specific database type. Then, physical design determines the storage and implementation details. Finally, the system is implemented and data is loaded. This structured process ensures a thorough and well-organized approach to creating a robust and efficient database system for users.

159. (A) The functional dependency is incorrect and violates BCNF.
In the given relation, the primary key is the combination of student ID and course ID. The dependency of course ID on professor name is problematic because professor name is not a superkey, yet it determines course ID. This violates Boyce-Codd Normal Form, which requires every determinant in a non-trivial functional dependency to be a superkey. The design should be decomposed to correct this structural issue.

160. (D) a, b, c, and d
Database management systems provide many significant advantages over traditional file-based systems. These include controlling redundancy, providing robust security and authorization, enforcing complex integrity constraints, and offering reliable backup and recovery services. All the provided statements correctly highlight these key benefits, which make a DBMS an essential tool for managing large and complex collections of data in a modern, secure, and highly efficient manner.

161. (B) DML
The update statement in SQL is used to modify existing data in a table, which is a core function of the Data Manipulation Language. DML commands are specifically designed for interacting with and managing the information stored within database objects. Unlike DDL for structure or DCL for security, DML provides the necessary tools for users to perform daily data maintenance tasks such as inserting and updating records.

162. (D) Statement I is incorrect, but Statement II is correct.
A view is a virtual table, not a physical one, meaning the first statement is incorrect. However, the second statement is correct because some views are updatable, and changes made to the data through the view can be reflected back into the underlying base tables. Understanding these properties is important for using views effectively to provide a flexible and secure interface to the actual database records.

163. (C) Unique and Not Null
A primary key is a fundamental constraint in a relational table that ensures each row can be uniquely identified. To achieve this, a primary key must satisfy two properties: its values must be unique across the entire table, and it cannot contain any null values. Together, these two requirements provide a reliable and consistent way to identify and reference individual records within the database system.

164. (D) Two or more transactions are waiting indefinitely for each other to release locks.
A deadlock is a concurrency problem that occurs when two or more transactions are stuck in a circular wait, each holding a lock that another needs and waiting for a lock that another holds. This situation prevents all involved transactions from making any progress. DBMSs use deadlock detection and resolution techniques, such as rolling back one of the transactions, to break these cycles and allow the system to proceed.

165. (D) a, b, c, d, and e
Temporal, multimedia, deductive, mobile, and geographic databases are all examples of enhanced data models designed for specialized requirements. These systems extend traditional relational capabilities to handle time, complex media, logical inference, mobile device challenges, and spatial data. All the provided statements accurately describe the unique features and focus areas of these advanced database technologies, reflecting the diverse and evolving landscape of modern data management.

166. (B) Star Schema
The star schema is a widely used data model in data warehousing, consisting of a large central fact table surrounded by several smaller dimension tables. This structure provides a simple and efficient way to organize data for analytical processing, allowing for fast queries and aggregations. The fact table contains the quantitative measures, while the dimension tables provide the descriptive context for the data being analyzed.

167. (D) Hidden Markov Model (HMM)
A Hidden Markov Model is a statistical tool used to model systems with hidden states that can be inferred from observable events. They are particularly effective for modeling sequential data

Downloaded from www.pdf.tube

without deleting the referencing rows entirely. By using these actions, designers can ensure that the database remains consistent even when related records are modified or removed from the parent table.

185. (A) A relational system must support a language that can handle data definition, manipulation, and transaction management.
Codd's Comprehensive Data Sublanguage Rule states that a truly relational system must support at least one language with a syntax that covers all aspects of database interaction. This includes defining the schema, manipulating the data, managing security, and controlling transactions. SQL is the most common example of such a comprehensive language, providing a single and consistent interface for all types of database management and user operations.

186. (B) `DELETE FROM table_name WHERE condition;`
To remove specific records from an SQL table, the delete from command is used along with a where clause to specify which rows should be deleted. This is a Data Manipulation Language command that removes data but leaves the table structure intact. Without a where clause, the command will remove all rows from the table, so it must be used carefully to avoid accidental data loss.

187. (C) The attributes of the relation have a specific order.
In the relational model, a relation is defined as an unordered set of tuples, and the attributes themselves also have no specific logical order. This means that the columns can be retrieved in any order without affecting the data. This property distinguishes relations from lists or arrays. All other options, such as atomic values and distinct tuples, are fundamental properties that every relation must strictly satisfy.

188. (D) Partial and transitive dependencies.
To reach Third Normal Form, a relation must already satisfy the requirements of 1NF and 2NF. This means it has already eliminated multi-valued attributes and any partial functional dependencies on the primary key. Reaching 3NF then requires the additional elimination of all transitive dependencies, where non-prime attributes depend on other non-prime attributes. Together, these steps result in a highly organized and redundant-free database table design.

189. (C) Committed
A transaction that has successfully finished its execution and has its changes permanently recorded in the database is said to be committed. Once committed, the changes are durable and will survive any subsequent system failure. The commit operation is a critical part of transaction management, signaling the successful completion of a logical unit of work and ensuring that its results are visible to other users.

NODIA

190. (D) a, b, c, and d
Query optimization is a critical part of modern database systems, aiming to find the most efficient way to execute queries. It uses execution plans, heuristic rules, and cost-based analysis to minimize execution time and resource usage. All the provided statements correctly describe different aspects of this process, highlighting its importance for high-performance data retrieval and the sophisticated techniques used by optimizers to achieve efficient results.

191. (C) Does not allow any locks to be released until the transaction commits or aborts.
The Strict Two-Phase Locking protocol is a more rigorous version of the basic 2PL. It requires that all locks held by a transaction, especially write locks, must be maintained until the transaction either commits or aborts. This prevents other transactions from reading uncommitted data, ensuring higher isolation and simplifying recovery. By delaying the release of locks, Strict 2PL provides a more robust and reliable method for managing concurrent operations.

192. (B) A system where the owner of a file can grant read/write access to other specific users.
Discretionary Access Control is a security model where the owner of an object has the authority to grant or revoke access permissions to other users. This is a flexible approach commonly used in many file systems and database environments. By allowing individual owners to manage their own data's security, DAC provides a decentralized way to control access based on user identity and specific permissions for different operations.

193. (A) Encapsulation

Encapsulation is a core principle of object-oriented programming that involves bundling data and the methods that operate on it into a single unit, or object. It also involves restricting direct access to some of the object's components, using public methods as the primary interface for interacting with the private data. This approach promotes modularity, data security, and simpler maintenance by hiding the internal implementation details from other parts.

194. (B) Genome Data Management
Genome data management is the most appropriate model for handling complex biological information such as gene sequences and protein structures. This field requires specialized database systems that are optimized for sequential data and can support sophisticated bioinformatics tools and analysis. These systems provide the necessary structure and performance for managing the massive and intricate datasets that are central to modern genomics and molecular biology research.

195. (B) a, b, c, and d only
OLTP and OLAP systems have very different focuses and structural characteristics. OLTP is transaction-oriented, deals with current data, is highly normalized, and handles simple queries. OLAP is analysis-oriented, uses historical data, is often denormalized, and performs complex analytical queries. The fifth statement is incorrect because OLTP is actually optimized for writes, while OLAP is optimized for reads. Therefore, the first four statements correctly compare the two systems.

196. (B) Association Rule Mining
The Apriori algorithm is a classic and widely used algorithm for association rule mining. It is designed to identify frequent itemsets in a transactional dataset and then use them to generate interesting association rules. Apriori works by iteratively identifying items that meet a minimum support threshold, making it a foundation for many applications in market basket analysis and other pattern discovery tasks in data mining.

197. (A) Sequencing Analysis
The pattern identifying that users who visit specific pages in a certain order are likely to make a purchase is an example of sequencing analysis. This technique focuses on discovering relationships between events that follow one another through time. By analyzing these sequences, businesses can gain valuable insights into user behavior and optimize their websites and marketing strategies to improve conversions and enhance the overall customer experience.

198. (B) The uncertainty and trustworthiness of data.
The veracity characteristic of Big Data refers to the inherent uncertainty and varying quality of the information being processed. Data from diverse sources like social media or sensors can be noisy, incomplete, or inaccurate. Managing this uncertainty is a major challenge in Big Data analysis, as researchers must evaluate the trustworthiness and reliability of the data to ensure that any derived insights are accurate and meaningful.

199. (C) YARN (Yet Another Resource Negotiator)
YARN is a core component of the Hadoop ecosystem that acts as the operating system for the cluster. Its primary role is to manage cluster resources and schedule jobs for different applications, including MapReduce. By decoupling resource management from data processing, YARN allows multiple applications to share the same cluster, improving efficiency and flexibility in large-scale Big Data environments where diverse tasks must be performed.

200. (A) Both A and R are true and R is the correct explanation of A.
Graph databases are exceptionally effective for managing data with many complex relationships because their model focuses directly on nodes and connections. This allows for extremely fast traversal of relationships without the need for the expensive join operations required by relational databases. The reason correctly explains how this direct focus on connectivity provides superior performance for applications like social networks, where understanding and querying relationships is a primary requirement.

201. (D) The application requires high scalability and a flexible data model for unstructured data.
A developer might choose a NoSQL database when an application needs to handle massive amounts of unstructured data or requires high

Click Here to Buy book on Amazon

horizontal scalability. NoSQL systems offer more flexible data models and can easily scale by adding more commodity servers, making them well-suited for modern, large-scale web applications. This distinguishes them from relational databases, which are better for well-defined, structured data and complex SQL join operations with strict ACID compliance.

202. (B) All of the above
A comprehensive Big Data architecture includes several essential components to handle the volume, velocity, and variety of information. These include diverse data sources for ingestion, specialized storage systems for managing massive datasets, and frameworks for both batch and real-time processing. By integrating all of these components, organizations can build a robust and scalable infrastructure for extracting valuable insights from their large and complex Big Data assets.

203. (C) Relational Algebra, Relational Calculus
Relational algebra is a procedural language where users specify the specific operations required to obtain a result. In contrast, relational calculus is a non-procedural language that focuses on describing what data is needed without specifying how to retrieve it. This pair represents the two fundamental types of formal query languages used in relational database theory, each providing a different and complementary way to interact with stored information.

204. (D) A relationship set
In an Entity-Relationship diagram, the diamond shape is the standard symbol used to represent a relationship set. These diamonds are linked to the participating entity sets, showing how they are connected and describing the associations between them. This visual notation is a key part of the ER model, providing a clear and consistent way to show the logical structure and relationships within a database design for users and developers.

205. (A) b, d, c, a
Normal forms provide a systematic way to organize database tables and minimize redundancy, progressing from less to more restrictive. 1NF requires atomic values, 2NF adds the elimination of partial dependencies, and 3NF removes transitive dependencies. BCNF is the strictest of these four, requiring that every determinant be a superkey. This sequence reflects the increasing rigor applied to ensure a highly structured and efficient database table design.

206. (B) A specific DML (INSERT, DELETE, UPDATE) event on a table.
A database trigger is a specialized procedure that executes automatically in response to specific data manipulation events on a table, such as an insert, delete, or update. Triggers are powerful tools for enforcing integrity constraints, maintaining audit logs, and automating related actions across multiple tables. By operating in response to DML events, triggers ensure that necessary business rules are consistently applied whenever the specified data modifications occur.

207. (C) Transitive dependency
In the relation R(A, B, C, D) with primary key AB, a transitive dependency exists because AB determines C, and C then determines D. Since both C and D are non-prime attributes, the dependency of D on AB is indirect, through C. This type of transitive dependency is what 3NF is specifically designed to identify and eliminate, ensuring that all non-prime attributes are dependent only on the primary key itself.

208. (A) Allow multiple transactions to execute simultaneously without causing inconsistencies.
The main goal of concurrency control in a database management system is to ensure that multiple users can execute transactions simultaneously without interfering with each other. This is essential for maintaining data consistency and accuracy in a multi-user environment. By using techniques like locking and timestamping, the system manages concurrent access and prevents problems like dirty reads or non-repeatable reads, ensuring reliable and accurate database operations.

209. (C) Object-oriented programming language capabilities.
An Object-Oriented Database Management System integrates traditional database capabilities with the features of object-oriented programming, such as complex objects, inheritance, and encapsulation. This allows developers to store and manage real-world objects more naturally

Downloaded from www.pdf.tube

within the database. By combining these two paradigms, OODBMSs provide a powerful and versatile framework for applications that require complex data modeling beyond the capabilities of the traditional and strictly relational database model.

210. (C) P-4, Q-2, R-1, S-3
Different data mining techniques are suited for various real-world use cases. Clustering is used for customer segmentation, classification for tasks like spam detection, and regression for predicting numerical values like house prices. Association rule mining is commonly used for market basket analysis. Matching these techniques to their typical applications provides a clear overview of how they are utilized to extract valuable knowledge and patterns from diverse datasets in many different fields.

211. (B) A collection of conceptual tools for describing data, data relationships, and consistency constraints.
A data model is a high-level conceptual framework that provides the tools necessary to describe the structure, relationships, and constraints of data within a database system. It serves as an abstraction that allows designers to focus on the logical organization of information without getting bogged down in physical storage details. Data models are essential for ensuring a consistent and well-organized approach to database design and management for users.

212. (A) Both Statement I and Statement II are correct.
In a two-tier architecture, tasks can be distributed in various ways between the client and the server. A three-tier architecture adds a specialized middle layer, usually an application server, to handle the business logic and act as an intermediary. Both statements accurately describe these two fundamental ways of organizing a database system's physical and logical components, highlighting the increasing complexity and flexibility of modern architectural designs.

213. (C) A relation containing all possible pairs of tuples where the first tuple is from R and the second is from S.
In relational algebra, the Cartesian product of two relations R and S results in a new relation that contains every possible combination of rows from the two input tables. If R has n rows and S has m rows, the product will have n times m rows. This operation is fundamental for combining data from different relations and serves as a basis for more complex operations like the join.

214. (D) The update on the referenced primary key also updates the corresponding foreign key values.
The on update cascade action in a foreign key constraint ensures that if a primary key value in a parent table is updated, all corresponding foreign key values in the child table are also automatically updated. This maintains referential integrity and ensures that the links between related records are kept accurate. It is a vital tool for ensuring that changes to identifying information are consistently reflected throughout the entire database.

215. (A) `GRANT`
The grant command is a key component of the Data Control Language in SQL. It is used by administrators to provide specific access privileges to users, such as the ability to read or modify data in certain tables. As a DCL command, grant is fundamental for managing database security and ensuring that users only have the permissions they need to perform their authorized tasks within the system.

216. (B) `STRUCTURE`
Standard SQL supports various data types like integer, decimal, and boolean to represent different kinds of information. However, structure is not a standard SQL data type. While some database systems may have similar concepts for complex data, it is not part of the basic set of types used in the standard SQL language. Therefore, it is the only incorrect option among the provided choices for valid SQL data types.

217. (D) All attributes are single-valued.
A database schema is in First Normal Form if every attribute contains only single, or atomic, values. This means that multi-valued attributes and nested groups are not permitted. By ensuring that each field has only one value from the domain, 1NF provides a consistent and simple structure that is essential for all subsequent normalization steps. Reaching 1NF is the

foundation for creating a well-structured and efficient relational database design.

218. (A) Database recovery
Shadow paging is a technique used in database systems primarily for recovery and for ensuring transaction atomicity. It works by maintaining two sets of page tables, the current and the shadow, to manage updates. If a transaction fails, the system can simply revert to the shadow page table to restore the previous consistent state. This provides a reliable way to ensure that the database remains accurate after a system failure or crash.

219. (C) OLAP (Online Analytical Processing)
Online Analytical Processing is a technology used for the multidimensional analysis of large datasets from different perspectives. OLAP systems allow users to quickly perform complex aggregations and summarizations, often using a data cube model. This makes OLAP an essential tool for data analysis, business intelligence, and reporting, providing insights that are not easily obtainable from traditional transactional systems. Its focus is on analysis rather than daily transaction processing.

220. (B) a, b, and d only
K-Nearest Neighbour, Support Vector Machine, and Hidden Markov Model are all recognized and widely used data mining techniques for classification, regression, and sequence modeling. Relational algebra, however, is a formal query language used to interact with relational databases, not a data mining technique for finding patterns. Therefore, the first, second, and fourth options represent the correct group of valid data mining methods used for knowledge discovery in large datasets.

221. (D) Map and Reduce
The MapReduce programming model is based on two primary functions: map and reduce. The map function processes input data to generate intermediate key-value pairs, while the reduce function then aggregates these pairs to produce the final result. By breaking down tasks into these two distinct phases, MapReduce provides a powerful and scalable way to process massive amounts of data in parallel across a distributed computing cluster.

222. (B) A move away from the relational model's dominance for certain types of applications.
The term NoSQL primarily signifies a shift away from the complete dominance of the relational model for all database applications. It highlights the emergence of alternative data models, such as document and graph, that are better suited for specific requirements like high scalability and handling unstructured data. NoSQL emphasizes that the traditional relational approach is not always the best choice for every modern data management challenge in distributed computing environments.

223. (A) `GRANT SELECT ON Employees TO John;`
To give a specific user the ability to query a table, the grant select command is the appropriate SQL statement. This Data Control Language command allows administrators to assign specific privileges to users on individual database objects. By granting select permission on the employees table to user John, the administrator ensures he can retrieve data from that table while following the established security policies for the database.

224. (B) A high-level, implementation-independent view of the entire database.
In the three-schema architecture, the conceptual schema provides a high-level, logical view of the entire database for the entire user community. It describes the data stored, the relationships between different entities, and any consistency constraints, while remaining independent of physical storage details and specific user views. This centralized logical description is essential for ensuring a consistent and well-organized approach to managing all the data within the system.

225. (D) a, b, c, d, and e
Relational database design involves several key concepts related to normalization and performance. Normalization aims to reduce redundancy and eliminate anomalies through stages like 2NF, 3NF, and BCNF. While BCNF is stricter than 3NF, denormalization is sometimes used to improve query performance at the cost of redundancy. All the provided statements correctly describe these fundamental principles and their roles in creating and maintaining an efficient and reliable database system for users.

Click Here to Buy book on Amazon

Key-value stores represent the simplest and most basic form of NoSQL databases, using a hash table model for data storage. In contrast, graph databases are specifically and highly specialized for managing data with a high degree of connectivity and complex relationships. Both statements accurately describe these two very different types of NoSQL systems, highlighting the wide diversity of data models available for various large-scale data management requirements.

243. (C) The effects of a transaction are not visible to other transactions until it is committed.
The isolation property of a database transaction guarantees that its intermediate effects are hidden from other concurrent transactions until it is successfully committed. This ensures that users do not read uncommitted, potentially inconsistent data. By maintaining this separation, isolation provides a consistent view of the database for each user, which is essential for maintaining the overall accuracy and integrity of information in a multi-user, transactional database environment.

244. (D) 3NF
In the given university table, a transitive dependency exists because course ID determines instructor ID, which in turn determines instructor name. Since course ID is the primary key and both instructor ID and name are non-prime attributes, this dependency violates Third Normal Form. Reaching 3NF would require decomposing the table into two, ensuring that all non-prime attributes are directly and fully dependent only on the primary key itself.

245. (D) a, b, c, and d
SQL provides a wide range of powerful tools for database management and user interaction. These include delete for data and drop for structure, views for abstraction and security, triggers for automated actions, and alter table for dynamic schema modifications. All the provided statements correctly describe these various features and their roles, highlighting the versatility and comprehensive nature of the SQL language for building and managing robust and secure database systems.

246. (B) It does not have enough attributes to form a primary key on its own and depends on a strong entity.
A weak entity set is a characteristic of the ER model where an entity cannot be uniquely identified by its own attributes alone. Instead, it must rely on a relationship with a strong identifying entity to form its complete primary key. This dependency is a key part of its definition and is visually represented by a double-lined rectangle in the ER diagram, helping designers identify entities with no independent identity.

247. (D) Classification is supervised learning (predefined classes), while clustering is unsupervised learning (no predefined classes).
The main difference between classification and clustering in data mining lies in the presence of predefined labels. Classification is a supervised learning task where the system is trained on labeled data to categorize new instances into predefined classes. Clustering is an unsupervised task that identifies natural groups in data based on similarities without using any prior knowledge or labels. This fundamental distinction defines the different approaches and applications for these two important data mining techniques.

248. (B) Three-tier architecture
A three-tier architecture is a common way to organize database systems, featuring a middle layer, typically an application server, between the user's client and the database server. This layer handles the business logic and acts as an intermediary, improving security, scalability, and maintainability. By separating the user interface, application processing, and data management, the three-tier model provides a flexible and robust framework for modern, large-scale database-driven applications for organizations.

249. (A) Its massive volume, high velocity, and wide variety.
Big Data is characterized by its massive volume, high velocity of generation and processing, and wide variety of data formats and sources. These three Vs represent the primary challenges that make Big Data difficult to manage with traditional relational database management systems. Special architectures and tools like Hadoop and NoSQL are necessary to handle these characteristics effectively and extract valuable insights from

234. (B) Null values must be treated consistently, primarily as unknown or inapplicable, and distinct from zero or empty strings.
Codd's rule for the systematic treatment of null values requires that they be handled consistently across the entire database management system. Nulls represent missing or inapplicable information and must be clearly distinguished from regular values like zero or an empty string. This consistent treatment is essential for ensuring accurate data analysis and reporting, as it provides a standardized way to represent and handle unknown or unavailable information in any context.

235. (A) `TRUNCATE TABLE Sales;`
The truncate table command is a fast and efficient way to remove all rows from a table without deleting the table structure itself. Unlike the delete command, which removes rows one by one and logs each operation, truncate is a structural operation that deallocates the data pages, making it much faster for large tables. This makes it an ideal tool for quickly clearing out a table's data while keeping it ready for new entries.

236. (B) Y is a subset of X.
A functional dependency X to Y is considered trivial if the dependent set of attributes Y is a subset of the determinant set X. This means that the dependency is naturally satisfied by the definition of the attributes and does not provide any additional information about the logical structure of the data. Identifying trivial dependencies is a basic step in the process of database normalization and understanding the functional relationships within a relation.

237. (C) Assigning a unique timestamp to each transaction and ordering operations based on these timestamps.
A timestamp-based protocol for concurrency control works by assigning a unique, monotonically increasing timestamp to each transaction as it begins. The system then uses these timestamps to order the execution of concurrent operations, ensuring that the final result is equivalent to a serial execution in timestamp order. This approach provides a reliable way to manage simultaneous transactions without the need for complex locking mechanisms or validation phases.

238. (D) Deductive Database
A deductive database is the most appropriate choice for building a system that can infer new facts from a set of rules and stored information. By combining a knowledge base of facts with a collection of logical rules, the system can automatically deduce information that is not directly present. This capability is essential for applications involving complex logical reasoning, such as identifying connections in a social network or answering sophisticated expert system queries.

239. (C) Fact and Dimension tables
Creating a data model for a data warehouse typically involves designing two types of tables: fact tables and dimension tables. Fact tables contain the quantitative measures of business processes, while dimension tables provide the descriptive context for that data. This star schema structure is the foundation of many data warehousing models, providing a simple and efficient way to organize and analyze large amounts of information from many different business perspectives.

240. (A) P-3, Q-4, R-2, S-1
HDFS provides distributed file storage, while MapReduce is a framework for parallel processing across a cluster. MongoDB is a popular NoSQL database, and OLAP uses a data cube model for multidimensional analysis. Matching these terms to their corresponding technologies provides a clear overview of the key components and concepts in the modern landscape of Big Data storage, processing, and analytical systems for large-scale information management.

241. (B) To store and retrieve data blocks as instructed by the NameNode.
In the Hadoop Distributed File System, the DataNodes are the workers of the cluster. Their primary role is to store and retrieve the actual data blocks as instructed by the master NameNode. They also handle the creation, deletion, and replication of these blocks across the cluster. Without the DataNodes, the system would have no place to store the massive amounts of data that are characteristic of modern Big Data applications.

242. (A) Both Statement I and Statement II are correct.

Downloaded from www.pdf.tube

226. (B) To serve as a basis for evaluating whether a DBMS can be considered truly relational.
Codd's 12 Rules were developed to provide a rigorous benchmark for determining whether a database management system can be considered truly relational. These rules define the fundamental requirements for data representation, manipulation, and integrity in a relational environment. By serving as a standard, Codd's rules have played a vital role in shaping the development and evaluation of modern relational database technology and its associated query languages.

227. (D) Tricking the application into including malicious SQL code in a database query.
An SQL injection attack works by manipulating user input to trick an application into including malicious SQL code within a database query. This can allow attackers to bypass security measures, steal sensitive data, or even modify or delete records. Protecting against SQL injection is a critical security requirement, often achieved through techniques like using parameterized queries and validating all user-provided data before it is used in a database operation.

228. (C) Interference between concurrently executing transactions.
The isolation property of ACID transactions is specifically designed to prevent interference between concurrently executing operations. It ensures that each transaction appears to run in isolation from others, even when they are accessing the same data simultaneously. By maintaining this separation, isolation prevents concurrency problems like dirty reads and non-repeatable reads, ensuring that the database remains in a consistent and accurate state despite the simultaneous actions of multiple users.

229. (A) Temporal Database
A temporal database is designed to store and manage data associated with time, making it suitable for historical stock market records. It allows efficient querying based on specific time ranges, tracking changes over periods, and maintaining time-stamped data. Since stock analysis often relies on past trends and timelines, temporal databases are more appropriate than multimedia, deductive, or mobile databases.

230. (A) Both A and R are true and R is the correct explanation of A.
The MapReduce programming model is inherently parallel because its map and reduce tasks can be executed independently on different subsets of data. This allows for massive speedups when processing large datasets across a distributed cluster of nodes. The reason correctly explains how this independent execution of tasks provides the parallel processing capability described in the assertion, making MapReduce a cornerstone of modern Big Data processing architectures.

231. (A) Its flexible schema, allowing documents in the same collection to have different structures.
The primary advantage of a document store NoSQL database like MongoDB is its flexible, schema-less design. This allows documents in the same collection to have different fields and structures, making it much easier to manage diverse and rapidly evolving data than the rigid tabular structure of a relational database. This flexibility is essential for modern web applications where data requirements can change frequently and information often arrives in various formats.

NODIA

232. (B) The processing load is distributed between a central server and multiple client machines.
A key characteristic of the client-server architecture is the distribution of the processing load between a central server and multiple client machines. The server typically handles data management and storage, while the clients handle the user interface and some application logic. This division of tasks provides several benefits, including better performance through distributed processing and more efficient use of resources across the entire network and database system.

233. (D) Join
In relational algebra, the join operation is used to combine tuples from two different relations based on a common attribute. This allows users to link related data from multiple tables together in a single result set. The join is one of the most powerful and frequently used operations in relational systems, providing the primary mechanism for integrating information that is distributed across different parts of the database schema.

large and complex datasets in diverse modern environments.

250. (C) a, b, c, and d only
NoSQL databases are specifically designed to offer horizontal scalability, making them highly effective for large-scale web applications and Big Data environments. These systems typically utilize flexible or non-existent schemas to accommodate various structured, semi-structured, and unstructured data types. Unlike traditional relational models, they encompass a wide variety of distinct data storage and retrieval approaches, meaning they are not a single type of database with only one model.

251. (B) The conceptual schema, which hides physical implementation details.
The conceptual schema provides a unified, logical representation of the entire database for the community of users. It effectively hides physical storage details and focuses on the logical structure of the data. This separation allows users to interact with a consistent data model without needing to understand the underlying hardware or physical file organization structures.

252. (A) Physical data independence.
Physical data independence is the property that permits modifications to the internal schema without altering the conceptual schema or application programs. This means changes in disk storage, file structures, or hardware do not force a rewrite of existing software. It ensures the database remains flexible and can evolve technologically while maintaining stable interfaces for users.

253. (C) To present a customized and relevant portion of the database to a user group.
The external schema, often referred to as the view level, describes a specific portion of the database relevant to a particular user group. By concealing the rest of the database, it simplifies the user experience and provides security. This layer ensures that different departments or roles only see and interact with data necessary for their tasks.

254. (D) facilitate flexibility and evolution of the database system over time.
Separating schemas within the three-schema architecture is crucial for achieving data independence. This structure allows the database system to evolve over time by facilitating changes at one level without impacting others. Consequently, the organization gains flexibility in upgrading hardware or modifying logical relationships while keeping the external user views and applications functional and stable throughout updates.

255. (C) These languages function as interfaces to define and manipulate the structures described by the schemas.
Data Definition Language and Data Manipulation Language serve as critical interfaces for managing database systems. While the definition language establishes the structure and rules of the schemas, the manipulation language provides the necessary tools for interacting with the actual data. Together, they allow users and administrators to effectively govern both the framework and the content of databases.

256. (B) That every row in a table must have a unique, non-null primary key.
Entity integrity mandates that every table must have a primary key to ensure each row can be uniquely identified. This constraint requires the primary key column to be both unique and non-null across the entire relation. Such rules prevent duplicate records and ensure that every entry in the database can be distinctly addressed and reliably retrieved by users.

257. (D) referential integrity.
Referential integrity governs the relationships between tables through the use of foreign keys. This constraint ensures that a value in a foreign key column must either match a valid primary key value in the referenced table or be null. By maintaining these links, the database prevents orphaned records and ensures consistent logical connections between different data entities.

258. (A) Relational Algebra is procedural, defining a sequence of operations, while Relational Calculus is non-procedural.
Relational algebra is a procedural language that specifies the exact sequence of operations required to obtain a query result. In contrast, relational calculus is a non-procedural language

Click Here to Buy book on Amazon

that describes the desired output without detailing the steps to compute it. Understanding these differences is fundamental for database professionals when evaluating query execution plans and various high-level database query languages.

259. (C) establish and maintain a logical link between two tables.
A foreign key consists of a column or set of columns in one table that refers to the primary key of another table. Its primary function is to establish and maintain a logical link between these two tables. This mechanism supports referential integrity by ensuring that the relationships between data entities remain consistent and valid across the database.

260. (A) a benchmark for evaluating the relational fidelity of a DBMS.
Codd's rules provide a set of thirteen comprehensive guidelines that a database management system must satisfy to be considered truly relational. These rules serve as a critical benchmark for evaluating the relational fidelity and capabilities of a system. They ensure that the software adheres to the core principles and mathematical foundations of the relational model for managing data.

261. (D) An `ALTER` statement adding a new column to a table.
Data Definition Language statements are used to manage the structure and schema of the database. Commands like creating, altering, or dropping tables fall under this category as they modify the underlying architecture rather than the data itself. Adding a new column to an existing table is a structural change that defines how data will be organized and stored.

262. (B) View.
A view acts as a virtual table that is based on the result-set of a specific SQL statement. It simplifies complex queries and enhances security by restricting access to certain subsets of data. Users interact with a view as if it were a real table, although the data actually resides in the underlying base tables of the database.

263. (C) It is a procedure that executes automatically in response to a DML event.
A database trigger is a specialized stored procedure designed to execute automatically in response to specific data manipulation events. These events typically include operations such as inserting, updating, or deleting records within a table. Triggers are essential for enforcing complex business rules, maintaining audit trails, and ensuring data integrity without requiring manual invocation by the user.

264. (D) improve performance and encourage the reuse of SQL code.
Stored procedures and functions consist of pre-compiled collections of SQL statements stored directly in the database. Utilizing these objects improves overall system performance because the logic is compiled once and executed multiple times. Additionally, they promote code reuse and modularity, allowing developers to centralize complex business logic and call it by name from various applications or different scripts.

265. (A) inserting malicious SQL code into data inputs for execution.
SQL injection is a critical security vulnerability that occurs when malicious code is inserted into input strings and subsequently executed by the database. Attackers exploit this flaw to bypass authentication, access sensitive information, or even delete data. It highlights the importance of implementing robust input validation and using parameterized queries to protect the integrity and confidentiality of systems.

266. (B) reduce data redundancy and prevent modification anomalies.
Normalization is a systematic design process used to organize columns and tables in a relational database. Its primary goal is to minimize data redundancy and prevent undesirable characteristics such as insertion, update, and deletion anomalies. By following these principles, designers ensure that the database remains accurate and consistent, while also simplifying the process of updating and maintaining information.

267. (B) functional dependencies.
Functional dependencies serve as the fundamental guide for the normalization process in relational databases. They describe relationships where the value of one attribute uniquely determines the

NODIA

value of another attribute within a relation. By analyzing these dependencies, database designers can identify redundancies and structural flaws, allowing them to decompose tables into smaller, more efficient forms that maintain data integrity.

268. (A) concurrency control techniques.
Concurrency control techniques are essential mechanisms used to manage simultaneous transactions in a database environment. These techniques prevent multiple users from interfering with each other when accessing or modifying the same data at the same time. By ensuring that transactions remain isolated and consistent, the system maintains data integrity even when many different operations are occurring across the network concurrently.

269. (C) A normalized schema can provide clearer data retrieval paths for the query optimizer.
Normalization significantly impacts query performance by organizing data into a more structured and logical form. A well-normalized schema provides the query optimizer with clearly defined paths for retrieving data efficiently. While optimization algorithms determine the best execution strategy, a clean database structure assists these processes by reducing redundancy and ensuring that data relationships are easily navigated during complex queries.

270. (D) indivisible and must be completed as a whole.
A transaction represents a logical unit of work that is characterized by its indivisibility. This means that all operations within the transaction must be executed successfully in their entirety, or none of them should be applied to the database. This all-or-nothing approach is fundamental for maintaining atomicity and ensuring the database remains in a consistent state after every operation.

271. (C) Temporal Database.
Temporal databases are specifically designed to manage time-varying data by incorporating valid time or transaction time components. This allows the system to systematically handle information that changes over different periods, making it ideal for historical analysis and auditing. Such databases enable users to reconstruct past states of data and understand how records have evolved throughout the history of the system.

272. (B) Mobile Databases.
Mobile databases are lightweight systems developed to operate effectively on devices with limited hardware resources, such as smartphones or sensors. These databases are uniquely capable of functioning in disconnected environments where continuous network access is unavailable. They include essential synchronization capabilities to ensure that data remains consistent with a central server once a network connection is eventually re-established by the device.

273. (D) use logical rules to infer new facts from existing data.
A deductive database extends the traditional relational model by incorporating logic programming capabilities. This feature allows users to define sophisticated rules and make logical inferences based on stored facts. Unlike standard databases that only return explicitly stored data, deductive systems can derive new information through automated reasoning, providing a more powerful tool for complex decision-making and advanced knowledge management applications.

NODIA

274. (A) An XML database for its flexible format.
XML and Internet databases offer a flexible format specifically designed for handling semi-structured data that does not conform to a rigid schema. These systems are ideal for storing data with nested elements and varying structures, such as web documents. By using query languages like XQuery, developers can efficiently retrieve and manage complex, hierarchical information that would be difficult to model.

275. (C) Multimedia Databases.
Multimedia databases are engineered to store and manage large collections of unstructured objects, including images, audio, and video files. A primary concern for these systems is content-based retrieval, which allows users to search for data based on its internal features rather than just metadata. This technology is crucial for applications that require identifying specific visual or auditory elements within vast repositories.

276. (A) It is a non-volatile and time-variant collection of data.

293. (D) lack of a rigid schema and complex joins.
Query optimization in NOSQL systems is particularly challenging because these databases lack the rigid schemas and support for complex joins found in traditional relational systems. Without a fixed structure, the query optimizer has fewer predefined rules to follow when determining the most efficient way to retrieve data. Consequently, developers must rely more heavily on careful data modeling and effective indexing strategies.

294. (B) The effective indexing and ordering of data sets.
Effective indexing and ordering of datasets are identified as crucial factors for achieving high query performance in NOSQL databases. Because these systems often handle massive scales and lack traditional relational structures, proper indexing is necessary to avoid inefficient full-table scans. Without these optimizations, queries can become extremely slow, negatively impacting the responsiveness and overall efficiency of modern, high-traffic web applications.

295. (C) distributed, scale-out architecture.
There is a strong alignment between NOSQL databases and cloud computing due to their distributed, scale-out architecture. This design perfectly matches the elastic and on-demand nature of cloud infrastructure, allowing organizations to easily add or remove resources as needed. By distributing data across many servers, NOSQL systems provide the high availability and performance required for contemporary, globally distributed cloud applications.

296. (D) To define and grant access rights and privileges to users.
Authorization is a core security mechanism used to define and grant specific access rights or privileges to users and applications. It dictates what data a user is permitted to see and which actions, such as selecting or updating records, they can perform. This process ensures that sensitive information is protected and that users only interact with data necessary for their roles.

297. (A) Because data is stored across multiple physical locations.
Security challenges are magnified in distributed database systems primarily because the data is stored across multiple, geographically dispersed physical locations. Managing security in this environment requires enforcing consistent access controls and authentication protocols across the entire network. Ensuring that a user's privileges are correctly applied regardless of which server processes the request is essential for protecting data from unauthorized access.

298. (B) authorization.
The enforcement of access rights that prevents a user from viewing specific information, such as account details from another branch, is a direct application of authorization. This security process ensures that even authenticated users are restricted to the data they are officially privileged to access. By managing these rights, the system protects confidentiality and prevents unauthorized browsing of sensitive customer information.

299. (B) data is stored in different locations but is managed as a single logical database.
A distributed database is defined as a system where data is physically stored across several different locations, while the management system treats the entire collection as a single logical database. This architecture provides users with a unified view of the data, regardless of its physical placement. Robust protocols are required to maintain consistency and security across these various nodes.

300. (A) user privileges are enforced consistently across all data nodes.
In a distributed client-server architecture, a robust security protocol must ensure that user privileges are enforced consistently across all nodes in the network. This means that access controls must be applied uniformly, regardless of which specific data server is processing a request. Such consistency is vital for preventing security gaps and ensuring that unauthorized users cannot access sensitive data segments.

Downloaded from www.pdf.tube

the nearest data points in the feature space, the algorithm assigns the most common class to the new object. This simple yet effective approach relies on the proximity of items to make predictions.

285. (C) clustering, as it groups similar customers without prior labels.
Clustering is the most appropriate technique for segmenting customers into groups when no prior knowledge or labels exist for those groups. This unsupervised method analyzes banking habits to find natural similarities and group customers who behave similarly. By identifying these distinct clusters, a bank can develop targeted marketing strategies and personalized services based on the actual observed patterns in customer behavior.

286. (D) Velocity, referring to the high speed of data generation.
Velocity is one of the core characteristics of Big Data, referring to the high speed at which data is generated, collected, and processed. In modern digital environments, information flows constantly from various sources like sensors, social media, and transactional systems. Efficiently handling this rapid influx is a major challenge for Big Data architectures, requiring specialized tools for real-time analysis.

287. (C) processing large datasets in parallel across a distributed cluster.
The Map-Reduce programming model is a powerful paradigm designed for processing and generating large datasets using parallel, distributed algorithms on a cluster of computers. By breaking down tasks into map and reduce phases, it allows for efficient computation across many nodes simultaneously. This distributed approach is essential for handling the massive scale of information typical in Big Data applications and environments.

288. (B) To provide a reliable and distributed storage layer for Big Data.
The Hadoop Distributed File System serves as the primary storage component within the Hadoop ecosystem, providing a reliable and distributed layer for Big Data. It is designed to run on commodity hardware and offers high-throughput access to large application datasets. By spreading data across multiple nodes, it ensures that information is readily available for processing by parallel algorithms like Map-Reduce.

289. (D) By replicating data blocks across multiple nodes.
HDFS achieves fault tolerance primarily by replicating data blocks across multiple nodes in a cluster. This redundancy ensures that if a single node fails, the data remains accessible from other locations. This design allows the system to continue operating reliably even when using inexpensive, failure-prone commodity hardware, making it a cost-effective solution for storing massive amounts of critical information.

290. (A) perform aggregation or summarization of the mapped data.
In the Map-Reduce programming model, the Reduce job is responsible for performing the final aggregation or summarization of data that has been processed and filtered by the Map phase. This stage takes the intermediate results and combines them to produce a concise and meaningful output. It is crucial for generating summaries, counting occurrences, or calculating totals from very large datasets.

291. (C) scalability and flexibility for web-scale applications.
The primary motivation behind the development of NOSQL databases was to address the growing need for scalability and flexibility in web-scale applications. Traditional relational models often struggle with the massive volume and diverse forms of data found in modern cloud environments. NOSQL systems prioritize performance and a flexible schema, allowing developers to build applications that can easily grow and adapt.

292. (A) Using specific APIs or query languages tailored to the data model.
Data in NOSQL systems is typically queried using specific APIs or query languages that are tailored to the unique data model of each product. For instance, document-oriented databases might use JSON-based structures for data retrieval. These specialized interfaces deviate from standard SQL syntax to better support the specific performance goals and flexible data structures inherent in various NOSQL database management systems.

NODIA

(A) P-4, Q-1, R-2, S-3
(B) P-1, Q-4, R-3, S-2
(C) P-4, Q-2, R-1, S-3
(D) P-3, Q-1, R-4, S-2

7. A system is in a state where a set of blocked processes each holds a resource and is waiting to acquire a resource held by another process in the set. This situation is best described as:
(A) Starvation
(B) Thrashing
(C) A race condition
(D) Deadlock

8. In memory management, the phenomenon where the system spends more time paging than executing is known as:
(A) Thrashing
(B) Swapping
(C) Demand paging
(D) Fragmentation

9. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : RAID Level 1 offers data redundancy by mirroring data on two disks.
Reason (R) : RAID Level 1 improves performance by striping data across multiple disks, allowing parallel reads and writes.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

10. The method of file access where records can be read and written out of order, in any place, is known as:
(A) Direct access
(B) Sequential access
(C) Indexed access
(D) Spooled access

11. Which of the following is NOT a primary service provided by an operating system?
(A) Program execution
(B) I/O operations
(C) File-system manipulation
(D) Application software compilation

12. Consider the following components related to system security:
1. Access Matrix
2. Cryptography
3. User Authentication
4. Program Threats
5. Revocation of Access Rights

Which of these are mechanisms or tools used for implementing security defenses?
(A) 1, 2, and 4 only
(B) 1, 2, 3, and 5 only
(C) 3, 4, and 5 only
(D) 1, 2, 3, 4, and 5

13. A virtual machine that is installed on a host operating system, which then runs on the hardware, is an example of what type of virtualization?
(A) Type 0 hypervisor
(B) Type 1 (bare-metal) hypervisor
(C) Type 2 (hosted) hypervisor
(D) Para-virtualization

14. In the context of the Linux Operating System, what is the primary role of Kernel Modules?
(A) To provide a comprehensive graphical environment for the user to interact with the system.
(B) To manage complex user authentication protocols and system-wide security policies.
(C) To determine the specific order in which processes are executed by the central processor.
(D) To load hardware drivers and system functions into the kernel dynamically as needed.

15. What is the correct chronological order of the following steps during the system boot process?
1. The operating system kernel is loaded into memory.
2. The bootstrap loader is executed.
3. System services and daemons are started.
4. The BIOS/UEFI performs a Power-On Self-Test (POST).

Choose the correct answer from the options given below:
(A) 2, 4, 1, 3
(B) 4, 2, 1, 3
(C) 4, 1, 2, 3
(D) 2, 1, 4, 3

Click Here to Buy book on Amazon

CHAPTER 05

System Software and Operating System

1. What is the primary function of a loader in the context of system software?
(A) To translate source programs written in high-level languages into binary machine code.
(B) To transfer an executable file from secondary storage into main memory for execution.
(C) To integrate multiple independent object modules into a single executable binary file.
(D) To coordinate the scheduling of processor time across various active system processes.

2. Which of the following best defines a system call in an operating system?
(A) The standard programmatic method for a user process to request a kernel-level service.
(B) An explicit request from one active process to another for interprocess data exchange.
(C) A hardware-level signal generated by a peripheral device to notify the processor.
(D) A standard function call executed entirely within the user-mode application space.

3. Consider the following statements regarding system software:
1. An interpreter translates and executes a program line by line.
2. A compiler translates the entire source program into an object program at once.
3. Relocation is the process of adjusting addresses in a program to run correctly in a different memory location.
4. A macro is a shorthand for a sequence of instructions.

Which of the statements given above are correct?
(A) 1 and 2 only
(B) 1, 2, and 3 only
(C) 3 and 4 only
(D) 1, 2, 3, and 4

4. Read the following statements and choose the correct option.

Statement I : Peterson's solution is a classic software-based solution to the critical-section problem.
Statement II : Peterson's solution is applicable to any number of processes and does not suffer from busy waiting.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

NODIA

5. Which of the following is an INCORRECT statement about threads?
(A) Threads within the same process share the same address space.
(B) Context switching between threads is generally faster than between processes.
(C) Each thread has its own independent process control block (PCB).
(D) Threads are often referred to as lightweight processes.

6. Match List-I with List-II regarding CPU scheduling algorithms and their characteristics and select the correct answer from the options given below.
List I:(Algorithm)
P. FCFS (First-Come, First-Served)
Q. SJF (Shortest Job First)
R. Round Robin
S. Priority Scheduling
List II:(Characteristic)
1. Optimal in terms of average waiting time for a given set of processes.
2. Designed specifically for time-sharing systems.
3. Can lead to starvation of low-priority processes.
4. Can result in the convoy effect.

correct.

24. The C-SCAN disk scheduling algorithm is a variant of the SCAN algorithm. How does C-SCAN primarily differ from SCAN?
(A) It services requests in one direction and returns to start without servicing the return trip.
(B) It moves the disk arm to the nearest pending request regardless of the current direction.
(C) It starts at cylinder zero and moves toward the last cylinder to ensure full coverage.
(D) It services requests bidirectionally but prioritizes read operations over write tasks.

25. The conceptual model that describes protection in terms of a set of objects, a set of domains, and a set of access rights is known as the:
(A) Access Control List (ACL)
(B) Capability List
(C) Access Matrix
(D) Security Kernel

26. Which of the following is NOT a form of interprocess communication (IPC) in the Linux operating system?
(A) Kernel Modules
(B) Pipes
(C) Message Queues
(D) Semaphores

27. The feature in the Windows OS that allows multiple users to have active sessions on a single computer, with each user's desktop and applications running concurrently, is called:
(A) System Restore
(B) Hardware Abstraction Layer (HAL)
(C) Windows Subsystem for Linux (WSL)
(D) Fast User Switching

28. Which of the following is an INCORRECT pair regarding system software concepts?
(A) Compiler - Translates entire program at once
(B) Loader - Modifies a program to run from a new memory location
(C) Interpreter - Executes program line-by-line
(D) Linker - Binds multiple object files into a single executable

29. Which of the following statements about operating system services is correct?
(A) The command-line interface (CLI) is the only user interface provided by modern systems.
(B) Error detection and response are handled solely by application programs, not the kernel.
(C) The OS provides services for both users and the system, like resource allocation and accounting.
(D) Communication between processes is not a primary service managed by the operating system.

30. Consider the following statements related to process synchronization:
1. The critical-section problem aims to design a protocol that processes can use to cooperate.
2. A race condition occurs when several processes access and manipulate the same data concurrently, and the outcome depends on the particular order in which the access takes place.
3. A semaphore is a hardware-only solution to the critical-section problem.
4. Peterson's Solution is restricted to two processes.

Which of the statements given above are correct?
(A) 1, 2, and 4 only
(B) 1 and 3 only
(C) 2, 3, and 4 only
(D) 1, 2, 3, and 4

31. In the context of threading, a thread library like Pthreads provides an API for:
(A) Creating and managing threads at the user level.
(B) Directly manipulating kernel-level scheduling queues.
(C) Implementing hardware-level multithreading.
(D) Managing virtual memory for threads.

32. Which scheduling criterion measures the total time from the submission of a process to its completion?
(A) Waiting time
(B) Turnaround time
(C) Response time
(D) CPU utilization

Downloaded from www.pdf.tube

16. One of the key design principles of the Windows Operating System is the separation of core functionalities into different layers. The component responsible for interfacing directly with the hardware is known as the:
(A) Executive
(B) Hardware Abstraction Layer (HAL)
(C) Kernel
(D) System Services

17. Which of the following is a key design issue specifically for Distributed Systems?
(A) Maximizing CPU utilization
(B) Ensuring robustness and fault tolerance across a network
(C) Providing a command-line interpreter
(D) Managing contiguous memory allocation

18. Consider the following statements about process management:
1. The ready queue contains all processes residing in main memory, ready and waiting to execute.
2. A context switch involves saving the state of the current process and loading the saved state for the new process.
3. Inter process communication can be achieved through either shared memory or message passing.
4. A long-term scheduler selects processes from a pool and loads them into memory for execution.

Which of the statements given above are correct?
(A) 1 and 3 only
(B) 2 and 4 only
(C) 1, 2, and 4 only
(D) 1, 2, 3, and 4

19. In the Many-to-One multithreading model:
(A) Each user-level thread is mapped to a dedicated kernel thread.
(B) Multiple user-level threads are mapped to a smaller or equal number of kernel threads.
(C) Multiple user-level threads are mapped to a single kernel thread.
(D) The kernel is unaware of the existence of user-level threads.

20. A CPU scheduling algorithm that can lead to starvation for longer processes is:
(A) First-Come, First-Served (FCFS)
(B) Round Robin (RR)
(C) Multilevel Queue Scheduling
(D) Shortest Job First (SJF)

21. Match List-I with List-II concerning deadlock handling methods and select the correct answer from the options given below.
List I:(Method)
P. Deadlock Prevention
Q. Deadlock Avoidance
R. Deadlock Detection
S. Recovery from Deadlock
List II:(Technique/Concept)
1. Using a wait-for graph to find cycles.
2. Terminating a process or preempting resources.
3. Ensuring that at least one of the four necessary conditions cannot hold.
4. Using an algorithm like the Banker's algorithm to ensure the system never enters an unsafe state.
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-3, R-1, S-4

22. What is the primary purpose of the Translation Look-aside Buffer (TLB)?
(A) To buffer data being transferred between the CPU and I/O devices.
(B) To store recently used page table entries for faster address translation.
(C) To hold the page table for the currently running process.
(D) To track free memory frames in the system.

23. Read the following statements and choose the correct option.
Statement I : Contiguous memory allocation can suffer from external fragmentation.
Statement II : Paging is a memory management scheme that permits the physical address space of a process to be non-contiguous, thereby eliminating external fragmentation.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is

NODIA

5. Priority Scheduling

Which of these algorithms can potentially suffer from starvation?

(A) 1 and 4 only
(B) 2, 3, and 5 only
(C) 1, 2, and 3 only
(D) 2, 3, 4, and 5

44. Read the following statements and choose the correct option.

Statement I : The `wait()` (or P) operation on a semaphore decrements the semaphore value.

Statement II : The `signal()` (or V) operation on a semaphore increments the semaphore value.

(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

45. The One-to-One multithreading model provides more concurrency than the Many-to-One model primarily because:

(A) It requires significantly less memory to implement within the system.
(B) It allows only one thread to execute on the CPU at any given time.
(C) A blocking system call by one thread does not block the entire process.
(D) Thread creation and management are handled entirely in the user space.

46. What is the chronological sequence of events for handling an I/O request?

1. The kernel issues commands to the device controller.
2. The device driver loads the appropriate registers in the device controller.
3. A user process issues a system call.
4. The device controller generates an interrupt on completion.

Choose the correct answer from the options given below:

(A) 3, 2, 1, 4
(B) 1, 3, 2, 4
(C) 3, 1, 2, 4
(D) 2, 3, 1, 4

47. In the context of deadlock, the condition where a process must be holding at least one resource and waiting to acquire additional resources that are currently being held by other processes is known as:

(A) Hold and Wait
(B) Mutual Exclusion
(C) No Preemption
(D) Circular Wait

48. A page fault occurs when:

(A) The requested page is in the TLB.
(B) The process tries to access a page that is not currently in main memory.
(C) The requested page is in main memory but not in the cache.
(D) The page table contains an invalid address.

49. Which of the following is NOT a method for free-space management on a disk?

(A) Bit vector
(B) Indexed allocation
(C) Linked list
(D) Grouping

50. In cryptography, a method where the same key is used for both encryption and decryption is called:

(A) Symmetric cryptography
(B) Asymmetric cryptography
(C) Public-key cryptography
(D) Hashing

51. Match List-I with List-II regarding their roles in the operating system and select the correct answer from the options given below.

List I:(Component)

P. Kernel I/O Subsystem
Q. Directory Implementation
R. System Boot
S. Communication in Client-Server Systems

List II:(Associated Concept)

1. Bootstrap Loader
2. Sockets or RPC
3. Buffering, Caching, Spooling
4. Linear list or Hash table

(A) P-4, Q-3, R-2, S-1
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-4, R-1, S-3

Click Here to Buy book on Amazon

33. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : Deadlock prevention is a set of methods for ensuring that at least one of the necessary conditions for deadlock cannot hold.
Reason (R) : One way to prevent the hold and wait condition is to require that a process request all its resources before it begins execution.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

34. The process of moving a process from memory to a backing store, and then bringing it back into memory for continued execution is called:
(A) Swapping
(B) Paging
(C) Thrashing
(D) Segmentation

35. The RAID level that provides block-level striping with distributed parity, offering a good balance of performance and redundancy, is:
(A) RAID 0
(B) RAID 1
(C) RAID 5
(D) RAID 10 (1+0)

36. A file system structure where directories can contain files and other directories is known as a:
(A) Single-level directory structure
(B) Two-level directory structure
(C) Acyclic-graph directory structure
(D) Tree-structured directory structure

37. A malicious program that hides within or looks like a legitimate program is a:
(A) Worm
(B) Virus
(C) Trojan horse
(D) Logic bomb

38. Which of the following is a key characteristic of a Type 1 (bare-metal) hypervisor?
(A) It runs as an application layer on top of a standard host operating system.
(B) It runs directly on the hardware to control it and manage guest systems.
(C) It provides virtualization services for applications rather than full systems.
(D) It requires the guest operating system to be modified before it can run.

39. In the context of the Linux kernel design, what is a primary advantage of its monolithic structure combined with modules?
(A) It is inherently more secure than a standard microkernel architectural design.
(B) It reduces the kernel memory footprint to the absolute smallest size possible.
(C) It ensures all drivers are statically linked into the kernel during compilation.
(D) It offers high performance for kernel communication while maintaining flexibility.

40. Which of the following statements about Network based Operating Systems is INCORRECT?
(A) They are designed to provide a high degree of transparency and fault tolerance.
(B) Users are aware of the existence of multiple machines.
(C) Each machine runs its own local operating system.
(D) Access to remote resources is a key feature.

41. Which of the following is a low-level programming language that uses mnemonics to represent basic machine instructions?
(A) High-Level Language
(B) Machine Language
(C) Assembly Language
(D) Natural Language

42. What is the role of the `fork()` system call in process management?
(A) It creates a new process that is a duplicate of the calling process.
(B) It terminates the currently running process.
(C) It replaces the current process's memory space with a new program.
(D) It waits for a child process to terminate.

43. Consider the following scheduling algorithms:
1. FCFS
2. Preemptive SJF
3. Non-preemptive SJF
4. Round Robin

Downloaded from www.pdf.tube

A data warehouse is defined as a non-volatile and time-variant collection of data used to support management's decision-making. Being non-volatile means that data is not frequently changed once it is entered, while time-variant indicates that the system maintains a historical perspective. These characteristics distinguish data warehouses from operational databases, which focus on current records and frequent transactional updates.

277. (D) OLTP involves frequent data modifications, while OLAP involves complex queries on historical data.
Online Transaction Processing systems are designed to handle frequent, short transactions involving constant data modifications like insertions and updates. In contrast, Online Analytical Processing systems are optimized for executing complex queries on large volumes of historical data. This distinction ensures that operational tasks remain fast and efficient, while analytical processes can provide deep insights for business intelligence and long-term strategic planning.

278. (B) Roll-up.
The roll-up operation is a fundamental OLAP technique that involves aggregating data to a higher level of abstraction within a concept hierarchy. For example, moving from daily sales records to a monthly summary demonstrates this process. It allows analysts to view generalized trends by summarizing detailed data points into broader categories, facilitating a more comprehensive understanding of performance over time.

279. (C) To allow data analysis at different levels of abstraction.
Concept hierarchies are essential components in data warehousing that allow users to view and analyze information at different levels of abstraction. By organizing data into logical layers, such as cities within states or months within years, these structures enable powerful analytical operations. This flexibility allows businesses to explore their data from broad summaries down to specific, granular details as needed.

280. (D) major subjects of the enterprise, such as customers or products.
The term subject-oriented describes how a data warehouse organizes information around major subjects of an enterprise, such as customers, products, or sales. Unlike operational databases that focus on specific business processes like order entry or invoicing, a subject-oriented approach provides a consolidated view of all relevant data for a particular area. This organization is vital for effective decision support.

281. (A) Association rule mining.
Association rule mining is a data mining technique used to discover interesting relationships and correlations between variables in large datasets. A classic application is market basket analysis, which identifies products that are frequently purchased together, such as bread and milk. These insights help retailers optimize store layouts, design effective promotional strategies, and improve overall sales by understanding customer purchasing patterns more deeply.

282. (B) Classification is a supervised technique using labeled data, while clustering is unsupervised.
Classification is a supervised learning technique that uses a training set of labeled data to build a model for predicting categories of new items. Clustering is an unsupervised technique that groups objects based on similarities without using predefined labels. This distinction is fundamental because classification requires known outcomes to learn from, while clustering discovers hidden patterns and natural groupings within data.

283. (B) Regression.
Regression analysis is a predictive data mining technique specifically used for estimating continuous numerical values, such as the selling price of a house. Unlike classification, which assigns items to discrete categories, regression models the relationship between dependent and independent variables. This approach is ideal for tasks where the goal is to forecast a precise amount based on various input features and characteristics.

284. (A) the class that a majority of its neighbors belong to.
The K-Nearest Neighbour algorithm is an instance-based learning method where an object is classified according to a majority vote of its closest neighbors. By analyzing the labels of

NODIA

52. What is the primary function of a debugger?
(A) To optimize the source code for faster execution.
(B) To translate assembly language into machine code.
(C) To help a programmer find errors (bugs) in a program.
(D) To load a program into memory for execution.

53. An operating system design that structures the OS by removing all nonessential components from the kernel and implementing them as user-level programs is a:
(A) Layered approach
(B) Monolithic approach
(C) Microkernel approach
(D) Modular approach

54. Consider the following statements regarding threads and processes:
1. Threads are more lightweight to create than processes.
2. Processes have separate address spaces, whereas threads share the address space of their parent process.
3. Communication between threads of the same process is typically faster than communication between two separate processes.
4. A change to a global variable by one thread is not visible to other threads in the same process.

Which of the statements given above are correct?
(A) 1 and 2 only
(B) 1, 2, and 3 only
(C) 3 and 4 only
(D) 1, 2, 3, and 4

55. In real-time CPU scheduling, a task that must meet its deadline, otherwise it may cause a catastrophic failure, is known as a:
(A) Soft real-time task
(B) Firm real-time task
(C) Hard real-time task
(D) Periodic task

56. In the Banker's algorithm for deadlock avoidance, a state is considered safe if:
(A) There are enough resources for all processes to run to completion immediately.
(B) The system can allocate resources in some order and still avoid a deadlock.
(C) No process in the entire system is currently holding any hardware resource.
(D) All processes have declared their maximum resource needs to the OS in advance.

57. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : Paging suffers from internal fragmentation.
Reason (R) : In a paged system, memory is allocated in fixed-size blocks called frames, and the last page of a process often does not fill the entire frame, leading to wasted space within the allocated frame.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

58. What is the primary purpose of file system mounting?
(A) To create a new file system on a disk partition.
(B) To make a file system on a storage device accessible to the OS at a specific location in the directory tree.
(C) To check a file system for errors and inconsistencies.
(D) To delete all files from a file system.

59. The revocation of access rights is a complex problem in security. Which of the following is a common method for handling revocation?
(A) Immediate revocation by searching access lists or capability lists.
(B) Periodic revocation where rights are automatically revoked after a set time.
(C) Indirect revocation using pointers or indirection.
(D) All of the above.

60. In Windows OS, the set of system components that run in kernel mode and provide core OS services (like I/O management, process management, security) is collectively known as the:
(A) Hardware Abstraction Layer (HAL)

Click Here to Buy book on Amazon

(B) Windows API
(C) Executive
(D) User-mode services

61. Which type of language requires a compiler or interpreter to be executed by a computer?
(A) Machine Language
(B) Assembly Language
(C) High-Level Language
(D) Binary Code

62. What is the process of adjusting program addresses to match the final memory location where the program will be loaded called?
(A) Linking
(B) Loading
(C) Compilation
(D) Relocation

63. Which of the following is an INCORRECT statement about operating system structure?
(A) A monolithic kernel has all OS services running in kernel space.
(B) A microkernel aims to have only essential functions in the kernel.
(C) A layered approach organizes the OS into a hierarchy of specific layers.
(D) A monolithic kernel is always slower than a microkernel due to its large size.

64. Communication in client-server systems can be achieved using various methods. Which method involves one process calling a procedure on a remote machine as if it were a local call?
(A) Sockets
(B) Remote Procedure Calls (RPC)
(C) Message Queues
(D) Pipes

65. Consider the following regarding thread issues:
1. `fork()` and `exec()` system calls
2. Signal handling
3. Thread cancellation
4. Thread-local storage
5. Scheduler activations

Which of the above are considered significant issues or topics in the context of designing and using multithreaded programs?
(A) 1 and 2 only
(B) 1, 2, and 3 only
(C) 3, 4, and 5 only
(D) 1, 2, 3, 4, and 5

66. Which scheduling algorithm is most suitable for a time-sharing interactive system?
(A) Shortest Job First (SJF)
(B) First-Come, First-Served (FCFS)
(C) Round Robin (RR)
(D) Priority Scheduling (non-preemptive)

67. To prevent the circular wait condition for deadlock, one can:
(A) Allow processes to hold resources while requesting more.
(B) Allow the system to preempt resources from processes.
(C) Enforce a total ordering of all resource types and require that each process requests resources in an increasing order.
(D) Ensure that a resource can be used by only one process at a time.

68. The page-replacement algorithm that replaces the page that will not be used for the longest period of time is known as:
(A) First-In, First-Out (FIFO)
(B) Least Recently Used (LRU)
(C) Optimal (OPT)
(D) Most Frequently Used (MFU)

69. The structure that organizes a mass-storage device into concentric rings is called:
(A) A sector
(B) A track
(C) A cylinder
(D) A platter

70. Which of the following is an example of a system threat, as opposed to a program threat?
(A) Trojan Horse
(B) Trap Door
(C) Worm
(D) Logic Bomb

NODIA

71. Match List-I with List-II concerning file allocation methods and select the correct answer from the options given below.
List I:(Allocation Method)
P. Contiguous Allocation
Q. Linked Allocation
R. Indexed Allocation
S. FAT (File Allocation Table)

List II:(Characteristic)
1. Suffers from external fragmentation.
2. A variation of linked allocation used by MS-DOS.
3. Each file has its own index block containing pointers to data blocks.
4. Each block contains a pointer to the next block in the file.

(A) P-1, Q-4, R-3, S-2
(B) P-4, Q-1, R-2, S-3
(C) P-1, Q-3, R-4, S-2
(D) P-2, Q-4, R-1, S-3

72. The concept of using a portion of the hard disk as an extension of RAM is known as:
(A) Virtualization
(B) Virtual Memory
(C) Cache memory
(D) Memory-mapped file

73. The primary design principle of the Linux operating system kernel is best described as:
(A) A pure microkernel architecture.
(B) A monolithic kernel with modular capabilities.
(C) A layered system similar to the original THE system.
(D) A hybrid kernel with extensive use of a Hardware Abstraction Layer.

74. Read the following statements and choose the correct option.
Statement I : A distributed operating system makes a network of computers appear as a single, powerful uniprocessor system to the user.
Statement II : A network operating system provides less transparency than a distributed operating system.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

75. What is the main purpose of a macro preprocessor?
(A) To convert the entire program into machine code.
(B) To expand macro definitions by replacing macro calls with the macro's body.
(C) To link object files and resolve external references.
(D) To detect syntax errors in the source code.

76. Arrange the following memory management techniques in chronological order of their typical historical development and adoption.
1. Paging
2. Swapping
3. Segmentation
4. Contiguous allocation with fixed partitions (MFT)

Choose the correct answer from the options given below:
(A) 4, 2, 3, 1
(B) 2, 4, 1, 3
(C) 1, 2, 3, 4
(D) 4, 3, 2, 1

77. In process scheduling, the transition of a process from the 'running' state to the 'ready' state is caused by:
(A) I/O or event wait
(B) An interrupt or scheduler dispatch
(C) I/O or event completion
(D) Termination

78. A semaphore used to provide mutual exclusion is often called a:
(A) Binary semaphore
(B) Counting semaphore
(C) Mutex lock
(D) Spinlock

79. Consider the following statements about CPU scheduling:
1. Turnaround time is the sum of the periods spent waiting in the ready queue, executing on the CPU, and doing I/O.
2. The convoy effect can occur in FCFS scheduling when a short process gets stuck behind a long process.
3. Preemptive scheduling allows a running process to be interrupted and moved to the ready queue.
4. Response time is the time from the submission of a request until the first response is produced.

Which of the statements given above are correct?

Downloaded from www.pdf.tube

(C) Many-to-Many
(D) Two-level model (a variation of Many-to-Many)

119. A CPU scheduling algorithm assigns a fixed time quantum to each process. If the process does not complete or block within this time, it is preempted and moved to the back of the ready queue. This algorithm is:
(A) Priority Scheduling
(B) Shortest Remaining Time First
(C) Round Robin
(D) First-Come, First-Served

120. Which of the following is an INCORRECT pair?
(A) Deadlock Prevention - Negating one of the four necessary conditions.
(B) Deadlock Avoidance - Using Banker's algorithm.
(C) Deadlock Detection - Periodically checking for cycles in a resource-allocation graph.
(D) Deadlock Recovery - Granting all resource requests immediately.

121. Match List-I with List-II regarding memory management concepts and select the correct answer from the options given below.
List I:(Concept)
P. Demand Paging
Q. Page Replacement
R. Thrashing
S. Segmentation
List II:(Description)
1. Logical address space is a collection of variable-sized segments.
2. Excessive paging activity leading to low CPU utilization.
3. Loading a page into memory only when it is needed.
4. Deciding which memory frame to free up when a new page must be brought in.
(A) P-1, Q-2, R-4, S-3
(B) P-3, Q-4, R-1, S-2
(C) P-4, Q-3, R-2, S-1
(D) P-3, Q-4, R-2, S-1

122. Memory-mapped files allow:
(A) A file on disk to be mapped to the virtual address space for routine memory access.
(B) THE memory to be mapped directly to the CPU's internal cache for faster processing.
(C) All files in a specific directory to be loaded into the main memory at once.
(D) A physical file system to be mapped onto another existing file system in the kernel.

NODIA

123. THE time it takes for the disk to rotate the desired sector to be under the read/write head is known as:
(A) Seek time
(B) Rotational latency
(C) Transfer time
(D) Access time

124. What is file sharing in the context of a multi-user file system?
(A) THE ability for multiple users to access the same file concurrently with appropriate protection.
(B) THE process of copying a file from one user's directory to another's.
(C) THE process of mounting a file system.
(D) THE allocation of disk space for a file.

125. Which of the following is NOT a type of program threat?
(A) Trojan Horse
(B) Trap Door
(C) Denial of Service (DoS)
(D) Logic Bomb

126. Read the following statements and choose the correct option.
Statement I : THE Linux kernel is monolithic, meaning all its functionality is compiled into a single, large executable.
Statement II : THE Linux kernel's monolithic design is made flexible by the use of loadable kernel modules, which allow new functionality to be added at runtime.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

127. Which of the following is a primary design principle or component of the Windows NT family of operating systems?

(D) Second Chance

90. In the context of a resource-allocation graph, a cycle is:
(A) A sufficient but not necessary condition for a deadlock.
(B) A necessary and sufficient condition for a deadlock if there is only one instance of each resource type.
(C) Neither necessary nor sufficient for a deadlock.
(D) A necessary but not sufficient condition for a deadlock if there are multiple instances of a resource type.

91. Match List-I with List-II regarding OS concepts and select the correct answer from the options given below.
List I:(Concept)
P. System Boot
Q. Interprocess Communication
R. Process Synchronization
S. Critical-Section Problem
List II:(Associated Term)
1. Semaphore
2. Peterson's Solution
3. Message Passing
4. Bootstrap Loader
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-3, R-2, S-1
(C) P-4, Q-3, R-1, S-2
(D) P-3, Q-4, R-1, S-2

92. Which of the following is a task of the linker?
(A) Translating source code into object code.
(B) Resolving external symbols and combining multiple object files.
(C) Loading the executable file into memory.
(D) Executing the program line by line.

93. Read the following statements and choose the correct option.
Statement I : The user interface (UI) of an operating system can be command-line based (CLI), graphics-based (GUI), or batch-based.
Statement II : System calls provide the interface between a running program and the operating system.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

94. THE Many-to-Many multithreading model multiplexes many user-level threads to a smaller or equal number of kernel threads. What is a major benefit of this model?
(A) It is the simplest model to implement within the operating system kernel.
(B) It suffers from the problem that a blocking system call blocks the process.
(C) It allows for both concurrency and efficiency by combining multiple models.
(D) It does not require any support from the operating system kernel at all.

95. IN multiple-processor scheduling, what is symmetric multiprocessing (SMP)?
(A) Where one processor is the master and others are slaves.
(B) Where each processor is self-scheduling and may have its own ready queue.
(C) Where processes are permanently assigned to a specific processor.
(D) Where only one processor can run kernel code at a time.

96. Which of the following is NOT one of the four necessary conditions for deadlock to occur?
(A) Mutual Exclusion
(B) Hold and Wait
(C) Preemption
(C) Circular Wait

97. Consider the following techniques and concepts in memory management:
1. Swapping
2. Paging
3. Segmentation
4. Demand Paging
5. Thrashing
Which of these are fundamental memory allocation or management schemes?
(A) 1, 2, and 3 only
(B) 1, 2, 3, and 4 only
(C) 4 and 5 only
(D) 1, 2, 3, 4, and 5

Click Here to Buy book on Amazon

(A) 1 and 3 only
(B) 2, 3, and 4 only
(C) 1, 2, and 4 only
(D) 1, 2, 3, and 4

80. What is a race condition in the context of process synchronization?
(A) A situation where two processes are competing to enter the critical section.
(B) A situation where outcome depends on the order of execution of multiple threads.
(C) A situation where a low-priority process is indefinitely postponed by others.
(D) A situation where a process holds a resource and waits for another resource.

81. A system that uses thread pools for creating and managing threads is an example of:
(A) A multithreading model
(B) Implicit threading
(C) A thread library
(D) A threading issue

82. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : The Banker's algorithm for deadlock avoidance requires processes to declare their maximum resource needs in advance.
Reason (R) : Without knowing the maximum needs, the algorithm cannot determine if granting a resource request will lead to an unsafe state.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

83. THE problem of external fragmentation exists in which of the following memory allocation schemes?
(A) Paging
(B) Pure Segmentation
(C) Inverted Page Tables
(D) Demand Paging

84. What is the primary function of the Kernel I/O Subsystem?
(A) To provide a standardized interface between device drivers and the rest of the kernel.
(B) To directly control the physical hardware of I/O devices within the system architecture.
(C) To allow user applications to access I/O hardware without any kernel-level intervention.
(D) To manage the hierarchical structure of the file system residing on the storage disk.

85. Which of the following is NOT a valid user authentication method?
(A) Knowledge-based (e.g., passwords)
(B) Possession-based (e.g., smart cards)
(C) Biometric-based (e.g., fingerprints)
(D) Access-matrix-based (e.g., permission tables)

NODIA

86. What is the main advantage of virtualization?
(A) It reduces the physical hardware requirements for running multiple operating systems.
(B) It allows a single OS to be partitioned into multiple isolated user environments.
(C) It guarantees that applications will run faster than on native hardware.
(D) It eliminates the need for system software like compilers and linkers.

87. Which component of the Windows operating system architecture is analogous to the concept of system calls in a UNIX/Linux system?
(A) The Hardware Abstraction Layer (HAL)
(B) The set of services exposed by the Executive
(C) The registry
(D) The .NET Framework

88. What is a Distributed File System (DFS)?
(A) A file system where data is stored redundantly on a single computer.
(B) A system that allows clients to access files stored across several network servers.
(C) A file system that uses a tree structure for its local directories and files.
(D) A file system designed specifically for the Linux kernel operating system.

89. A system has a page-replacement policy where the page that has not been used for the longest amount of time is replaced. This policy is:
(A) Optimal (OPT)
(B) First-In, First-Out (FIFO)
(C) Least Recently Used (LRU)

(D) Kernel-level threads.

109. What is the primary purpose of user authentication in an operating system?
(A) To encrypt files stored on the disk.
(B) To ensure that a user is who they claim to be before granting access to the system.
(C) To monitor network traffic for malicious activity.
(D) To track resource usage for accounting purposes.

110. Which of the following describes the concept of virtualization?
(A) The process of compiling a program for a specific CPU architecture.
(B) The creation of a virtual (rather than actual) version of something, such as a hardware platform, operating system, or storage device.
(C) A security mechanism that isolates processes from each other.
(D) A technique for increasing disk access speed by reorganizing files.

111. Consider the following statements about the Windows file system, NTFS:
1. It supports features like file permissions, encryption, and compression.
2. It uses a Master File Table (MFT) to store information about files and directories.
3. It is based on the older FAT file system.
4. It does not support journaling, which can lead to data loss.

Which of the statements given above are correct?
(A) 1 and 2 only
(B) 2 and 3 only
(C) 1, 3 and 4 only
(D) 1, 2, 3, and 4

112. THE protocol that defines the structure of communication between components in a distributed system is known as the:
(A) Network Structure
(B) Communication Protocol
(C) Design Issue
(D) Robustness Protocol

113. What is the fundamental difference between a compiler and an interpreter?
(A) A compiler produces machine code, while an interpreter produces assembly code.
(B) A compiler is used for system software, while an interpreter is for applications.
(C) A compiler translates the entire program before execution, while an interpreter does so line by line.
(D) A compiler is faster, but an interpreter produces significantly more optimized code.

114. THE part of the operating system that permanently resides in main memory is commonly known as the:
(A) System Library
(B) Shell
(C) Kernel
(D) API

115. Arrange the following CPU scheduling algorithms from simplest/oldest to more complex/modern in their conceptual development.
1. Round Robin (RR)
2. First-Come, First-Served (FCFS)
3. Multilevel Feedback Queue Scheduling
4. Shortest Job First (SJF)

Choose the correct answer from the options given below:
(A) 2, 4, 1, 3
(B) 1, 2, 4, 3
(C) 2, 1, 4, 3
(D) 4, 2, 3, 1

NODIA

116. A situation where a process is ready to run but is waiting for the processor is found in the:
(A) Running state
(B) Ready queue
(C) Wait queue
(D) Job queue

117. Peterson's solution for the critical-section problem is based on which two ideas?
(A) Test-and-Set and Swap instructions
(B) Semaphores and Monitors
(C) Turn-taking and willingness to enter
(D) Message passing and Shared memory

118. Which multithreading model is most common in modern operating systems like Linux, macOS, and Windows?
(A) Many-to-One
(B) One-to-One

137. Which page replacement algorithm suffers from Belady's Anomaly?
(A) Least Recently Used (LRU)
(B) First-In, First-Out (FIFO)
(C) Optimal (OPT)
(D) Least Frequently Used (LFU)

138. THE part of the disk that passes under the head as it rotates is a:
(A) Cylinder
(B) Track
(C) Sector
(D) Platter

139. In the context of directory and disk structure, a directory can be implemented as:
(A) A sequential file containing file names and their pointers.
(B) A hash table for efficient lookup of file entries.
(C) Both a linear list and a hash table are possible implementations.
(D) Only as a special type of hardware register.

140. Which of the following is NOT a goal of protection in a computer system?
(A) Preventing processes from interfering with each other.
(B) Ensuring that each object is accessed correctly and only by those processes that are allowed to do so.
(C) Enforcing policies for resource usage.
(D) Maximizing the speed of I/O operations.

141. THE ability to add functionality to the Linux kernel at runtime is provided by:
(A) System Calls
(B) Kernel Modules
(C) THE BASH shell
(D) THE GNU C Compiler

142. A key design goal of a Distributed File System (DFS) is:
(A) Location transparency
(B) Maximizing disk fragmentation
(C) Ensuring all files are stored on a single central server
(D) Limiting access to local files only

143. Which of the following is an INCORRECT statement regarding system software?
(A) A macro is a unit of specification for program generation through expansion.
(B) A debugger is a tool used to test and find errors in other programs.
(C) Loading is the process of resolving references between different object files.
(D) An interpreter reads a source statement, determines its meaning, and executes it.

144. THE bootstrap program that starts the operating system is typically stored in:
(A) RAM
(B) THE hard disk's boot sector
(C) ROM or EEPROM (firmware)
(D) THE CPU cache

145. Consider the following statements about process synchronization:
1. THE critical-section problem requires a solution that satisfies mutual exclusion, progress, and bounded waiting.
2. A semaphore is a software-based synchronization tool.
3. Busy waiting is a characteristic of spinlocks.
4. Peterson's solution works for n-processes.

Which of the statements given above are correct?
(A) 1 and 2 only
(B) 1, 2, and 3 only
(C) 3 and 4 only
(D) 1, 2, 3, and 4

146. Which of the following is a primary disadvantage of the Many-to-One multithreading model?
(A) It requires significant kernel support for each created thread.
(B) Creating user threads is very slow compared to other models.
(C) A blocking system call by one thread blocks the entire process.
(D) It does not allow for true parallelism on a multicore system.

147. In a multiprocessor system, processor affinity refers to:
(A) THE tendency for a process to be assigned to the same processor it was previously running on.
(B) THE speed at which a processor can access main memory.
(C) THE number of cores available on a processor.

Downloaded from www.pdf.tube

(A) A pure monolithic kernel for maximum performance.
(B) A layered architecture with a Hardware Abstraction Layer (HAL) for portability.
(C) Primary reliance on a command-line interface with no native GUI support.
(D) An open-source development model.

128. Robustness in a distributed system refers to its ability to:
(A) Handle a large number of concurrent users.
(B) Continue operating correctly despite failures of some of its components.
(C) Provide a single-system image to the users.
(D) Secure communication channels between nodes.

129. Consider the following:
1. Assembly Language
2. Machine Language
3. High-Level Language

What is the correct order of abstraction from lowest (closest to hardware) to highest?
(A) 2, 1, 3
(B) 1, 2, 3
(C) 3, 1, 2
(D) 2, 3, 1

130. THE state of a process is defined in part by its:
(A) Current activity, represented by the value of the program counter.
(B) Parent process ID.
(C) Position in the job queue.
(D) Physical memory address.

131. Consider the following statements about process scheduling and operations:
1. THE long-term scheduler controls the degree of multiprogramming.
2. THE short-term scheduler (CPU scheduler) selects a process from the ready queue for execution.
3. A process is I/O-bound if it spends more time doing I/O than computations.
4. A context switch is a necessary part of interprocess communication.

Which of the statements given above are correct?
(A) 1 and 2 only
(B) 3 and 4 only
(C) 1, 2, and 3 only
(D) 1, 2, 3, and 4

132. THE mechanism that allows a thread to be notified by another thread that it should terminate, either immediately or at a convenient point, is called:
(A) Signal handling
(B) Thread cancellation
(C) Thread pooling
(D) Asynchronous procedure call

133. Which of the following is an example of implicit threading?
(A) Pthreads library
(B) Win32 threads API
(C) Java threads
(D) OpenMP

134. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : In a priority scheduling algorithm, a major problem is indefinite blocking, or starvation.
Reason (R) : Starvation can be solved by aging, which involves gradually increasing the priority of processes that wait in the system for a long time.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

135. Recovery from deadlock can involve process termination. Which of the following is a reasonable strategy for choosing a victim process to terminate?
(A) Terminate the process that has been running for the longest time.
(B) Terminate the process with the highest priority.
(C) Terminate the process that has consumed the least amount of CPU time.
(D) Terminate all deadlocked processes.

136. THE address generated by the CPU is referred to as a:
(A) Physical address
(B) Logical address
(C) Absolute address
(D) Relocated address

98. A RAID structure that provides mirroring without parity or striping is:
(A) RAID 0
(B) RAID 1
(C) RAID 5
(D) RAID 6

99. In file systems, what is the purpose of a bit vector in free-space management?
(A) To store pointers to all the free blocks on the disk.
(B) To represent disk space as a sequence of bits indicating block status.
(C) To link all the free blocks together in a list.
(D) To store the names of all the files in the directory.

100. A program that uses a flaw in the operating system to gain unauthorized access or privileges is exploiting a:
(A) System threat
(B) Trojan horse
(C) Network threat
(D) Program threat

101. What distinguishes a network operating system from a distributed operating system?
(A) A network OS provides a seamless image, while a distributed OS does not.
(B) A distributed OS is aware of the network, but users manage machines.
(C) A network OS is less transparent than a distributed operating system.
(D) There is no difference; the terms are considered entirely synonymous.

102. In the context of the Linux file system, the inode structure stores:
(A) The name of the file.
(B) The actual data content of the file.
(C) Metadata such as owner, permissions, and pointers to data blocks.
(D) The path of the file within the directory tree.

103. Which is the correct pair of a concept and its description?
(A) Process Scheduling - Selecting a process from the job pool on disk.
(B) Interprocess Communication - The kernel's ability to switch between processes.
(C) Semaphore - A variable used for signaling among processes.
(D) Critical Section - The part of the OS that handles interrupts.

104. Which of the following is NOT a CPU scheduling criterion?
(A) Throughput
(B) Waiting Time
(C) Disk Bandwidth
(D) Response Time

105. Read the following assertion and reason carefully and choose the correct option.
Assertion (A) : In deadlock detection, a wait-for graph can be used to detect a cycle among processes.
Reason (R) : A cycle in a single-instance resource wait-for graph indicates the existence of a deadlock.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

106. What is the difference between a page and a frame in memory management?
(A) A page is physical memory, while a frame is a block of logical memory.
(B) A page is logical memory, while a frame is a block of physical memory.
(C) Pages are used in segmentation, while frames are used in paging.
(D) There is no difference; the terms are considered to be interchangeable.

107. Which disk scheduling algorithm is designed to be fair and avoid starvation by servicing requests in both directions of head movement?
(A) FCFS
(B) SSTF
(C) SCAN (Elevator)
(D) C-LOOK

108. THE 'Application I/O Interface' standardizes the way applications interact with the OS for I/O. This is typically achieved through:
(A) A set of system calls.
(B) Direct hardware manipulation.
(C) A library of device drivers.

Click Here to Buy book on Amazon

(D) THE scheduling algorithm used by the processor.

148. Which of the following is a method for deadlock prevention?
(A) Use a wait-for graph to detect cycles.
(B) Use THE Banker's algorithm to ensure a safe state.
(C) Require processes to request resources in a predefined order.
(D) Terminate a process to break THE deadlock.

149. Match List-I with List-II regarding security threats and select the correct answer from the options given below.
List I:(Threat)
P. Trojan Horse
Q. Worm
R. Denial of Service
S. Trap Door
List II:(Description)
1. A program that uses networking to self-replicate and spread.
2. An attack designed to prevent a system from providing service.
3. A hidden segment of a program that allows for unauthorized access.
4. A malicious program disguised as a legitimate one.
(A) P-4, Q-1, R-2, S-3
(B) P-1, Q-4, R-3, S-2
(C) P-4, Q-2, R-1, S-3
(D) P-3, Q-1, R-4, S-2

150. A pure segmentation memory management scheme has which major drawback?
(A) It is limited to a small number of segments per process.
(B) It can suffer from significant external fragmentation.
(C) It requires a very large page table.
(D) It does not support dynamic linking.

151. Read the following statements and choose the correct option.
Statement I : Contiguous allocation of files on a disk leads to excellent read performance.
Statement II : Contiguous allocation of files suffers from THE problem of finding contiguous blocks of space for new files, leading to external fragmentation.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

152. Which RAID level is known as striping and offers no redundancy?
(A) RAID 0
(B) RAID 1
(C) RAID 5
(D) RAID 6

153. Spooling is a form of I/O management that:
(A) Uses interrupts to signal I/O completion.
(B) Allows THE CPU to poll a device for its status continuously.
(C) Holds data for a device in a buffer on disk until THE device is ready.
(D) Allows direct memory access for high-speed devices.

154. What does revocation of access rights mean in THE context of system security?
(A) Granting a user new permissions.
(B) Taking away previously granted permissions from a user or process.
(C) Verifying a user's identity.
(D) Encrypting a user's data.

155. Which of THE following is NOT a feature of THE NTFS file system used in Windows?
(A) Journaling
(B) Support for large files and volumes
(C) Access control lists (ACLs) for security
(D) A fixed, 8.3 filename limitation

156. In distributed systems, what is a primary purpose of a communication protocol like TCP/IP?
(A) To manage process scheduling on a single machine.
(B) To provide a standard way for machines on a network to exchange data.
(C) To allocate memory to processes.
(D) To implement a file system.

NODIA

214. Which of THE following statements about threading is INCORRECT?
(A) User-level threads are generally faster to create and manage than kernel-level threads.
(B) THE one-to-one model maps each user thread to a kernel thread.
(C) IN THE many-to-one model, if one thread performs a blocking system call, THE entire process is blocked.
(D) Kernel-level threads are invisible to THE operating system.

215. Consider a preemptive priority scheduling algorithm. A low-priority process may be indefinitely postponed from running if there is a steady stream of high-priority processes. This problem is known as:
(A) Deadlock
(B) Starvation
(C) Convoy effect
(D) Aging

216. Which of THE following is NOT a method for handling deadlocks?
(A) Prevention
(B) Avoidance
(C) Detection and Recovery
(D) Augmentation

217. IN a paging system, THE effective memory access time is influenced by:
(A) THE page fault rate and THE time to service a page fault.
(B) THE size of THE physical memory only.
(C) THE number of processes in THE ready queue.
(D) THE speed of THE CPU clock.

218. THE information about all files in a volume is kept in a structure known as a:
(A) File Control Block (FCB)
(B) Directory or Volume Table of Contents
(C) Boot block
(D) Superblock

219. Read THE following assertion and reason carefully and choose THE correct option.
Assertion (A) : Cryptography is a fundamental tool for implementing security in computer systems.
Reason (R) : It can provide confidentiality and integrity of data by transforming it into a non-readable format.
(A) Both A and R are true and R is THE correct explanation of A.
(B) Both A and R are true but R is not THE correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

220. Which type of virtual machine runs on a host operating system?
(A) Type 1 (bare-metal)
(B) Type 2 (hosted)
(C) Paravirtualized VM
(D) System VM

221. Match List-I with List-II concerning Linux OS concepts and select THE correct answer from THE options given below.
List I:(Linux Concept)
P. Process Management
Q. File System
R. Scheduling
S. Interprocess Communication
List II:(Associated Term/Feature)
1. Completely Fair Scheduler (CFS)
2. ext4
3. fork() and exec()
4. Pipes, sockets
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-2, R-4, S-1
(C) P-3, Q-2, R-1, S-4
(D) P-2, Q-3, R-1, S-4

222. Which of THE following is a key design principle for THE Windows operating system?
(A) Adherence to THE POSIX standard for maximum compatibility with UNIX.
(B) A client-server model for internal operations, with a microkernel-like base.
(C) A purely monolithic kernel structure.
(D) Open-source code and community-driven development.

223. A key characteristic of a Network based Operating System is that:
(A) It presents a single-system image to THE user, hiding THE network.
(B) Each computer on THE network runs its own OS and users are aware of different

Downloaded from www.pdf.tube

access.
(B) An object, specifying which users/processes have what access rights to it.
(C) A user, specifying their password and authentication details.
(D) THE kernel, specifying THE order of process execution.

168. In Linux process management, a call to `fork()` followed by a call to `exec()` is used to:
(A) Terminate THE current process.
(B) Create a new thread within THE current process.
(C) Create a child process and have it run a new program.
(D) Send a signal to another process.

169. What is a protocol in THE context of communication structure?
(A) A hardware device used for network connectivity.
(B) A set of rules that govern THE format and transmission of data.
(C) THE physical layout of a network (e.g., bus, star).
(D) A software application for sending emails.

170. Which of THE following is an example of a High-Level Language?
(A) Intel x86 assembly
(B) Binary machine code
(C) Java
(D) MIPS instruction set

171. A process control block (PCB) does NOT contain which of THE following?
(A) THE process state
(B) CPU registers
(C) THE contents of THE user stack
(D) A list of open files

172. Which of THE following statements about interprocess communication (IPC) is correct?
(A) Shared memory IPC is generally faster than message passing IPC.
(B) Message passing requires processes to share a common address space.
(C) Shared memory is less prone to synchronization errors than message passing.
(D) Sockets are a mechanism for IPC between processes on THE same machine only.

173. Arrange THE following steps in handling a page fault in a demand paging system chronologically.
1. THE operating system finds a free frame.
2. THE process is restarted at THE instruction that caused THE trap.
3. A trap to THE operating system is generated.
4. THE required page is loaded from disk into THE free frame.
5. THE page table is updated.

Choose THE correct answer from THE options given below:
(A) 3, 1, 4, 5, 2
(B) 1, 3, 4, 5, 2
(C) 3, 4, 1, 2, 5
(D) 4, 3, 1, 5, 2

174. Which of THE following is a key issue in multithreaded programming related to signal handling?
(A) Determining which thread in a process should receive a signal.
(B) Ensuring all signals are ignored by all threads.
(C) Converting signals into system calls automatically.
(D) Preventing signals from ever being generated.

175. What is THE primary objective of CPU scheduling?
(A) To keep THE CPU as busy as possible and provide fast response times to users.
(B) To ensure all processes get an equal amount of CPU time.
(C) To load processes from disk into main memory.
(D) To minimize THE number of context switches.

176. THE condition that specifies that a resource can be released only voluntarily by THE process holding it, after that process has completed its task, is known as:
(A) Mutual Exclusion
(B) Hold and Wait
(C) No Preemption
(D) Circular Wait

NODIA

177. Thrashing can be reduced or eliminated by:

implement directly.

Statement II : Access Control Lists (ACLs) and Capability Lists are two common implementations of THE Access Matrix concept.

(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

235. THE main function of THE Hardware Abstraction Layer (HAL) in THE Windows OS is to:

(A) Provide THE graphical user interface.
(B) Hide hardware differences from upper levels of THE system.
(C) Manage THE file system.
(D) Schedule threads for execution.

236. A software that translates an assembly language program into machine code is called an:

(A) Interpreter
(B) Compiler
(C) Assembler
(D) Linker

237. THE long-term scheduler is responsible for:

(A) Selecting which process should be executed next by THE CPU.
(B) Swapping processes in and out of main memory.
(C) Selecting processes from THE job pool on disk to be loaded into memory.
(D) Allocating I/O devices to processes.

238. Consider THE following statements regarding operating system security and threats:

1. A worm is a standalone program that replicates itself to spread to other computers.
2. A virus requires a host program to attach itself to.
3. Cryptography can be used to ensure user authentication.
4. An access matrix defines THE rights of each process with respect to each object.

Which of THE statements given above are correct?

(A) 1 and 2 only
(B) 1, 2, and 4 only
(C) 3 and 4 only
(D) 1, 2, 3, and 4

239. THE issue of how to handle signals in a multithreaded environment is complex. Which of THE following is a common option for signal handling?

(A) Deliver THE signal to THE thread to which THE signal applies.
(B) Deliver THE signal to every thread in THE process.
(C) Deliver THE signal to certain threads in THE process.
(D) All of THE above.

240. IN real-time CPU scheduling, 'soft' real-time systems are those where:

(A) A critical real-time task must be completed within its deadline.
(B) A critical real-time task gets priority over other tasks, but THE deadline is not mandatory.
(C) There are no real-time tasks.
(D) All tasks have THE same priority level.

241. Read THE following assertion and reason carefully and choose THE correct option.

Assertion (A) : IN a system with multiple instances of each resource type, a cycle in THE resource-allocation graph is a necessary but not a sufficient condition for deadlock.

Reason (R) : Even if a cycle exists, if a process outside THE cycle releases its resources, one of THE processes in THE cycle might be able to acquire THE needed resource and break THE cycle.

(A) Both A and R are true and R is THE correct explanation of A.
(B) Both A and R are true but R is not THE correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

242. What is Belady's Anomaly in THE context of page replacement algorithms?

(A) THE anomaly where THE page fault rate increases as THE number of allocated frames increases.
(B) THE anomaly where THE CPU utilization decreases as THE number of processes

Click Here to Buy book on Amazon

186. Consider THE following statements about process states:
1. 'New': THE process is being created.
2. 'Running': Instructions are being executed.
3. 'Waiting': THE process is waiting for some event to occur (such as an I/O completion).
4. 'Terminated': THE process has finished execution.
5. 'Ready': THE process is waiting to be assigned to a processor.

Which of these are standard states in a typical process life cycle model?
(A) 2, 3, and 5 only
(B) 1, 2, 3, and 4 only
(C) 1, 2, 3, 4, and 5
(D) 2 and 5 only

187. THE critical-section problem must satisfy which of THE following requirements?
(A) Mutual Exclusion, Progress, Bounded Waiting
(B) Preemption, Circular Wait, Hold and Wait
(C) Turnaround Time, Waiting Time, Response Time
(D) Atomicity, Consistency, Isolation, Durability

188. IN THE context of thread libraries, what does `Pthreads` refer to?
(A) A proprietary thread library for Windows systems.
(B) A standard IEEE specification for thread behavior.
(C) A hardware implementation of threading.
(D) A type of implicit threading.

189. Which scheduling algorithm gives THE minimum average waiting time for a given set of processes?
(A) FCFS
(B) SJF (Shortest Job First)
(C) Round Robin
(D) Priority Scheduling

190. Read THE following assertion and reason carefully and choose THE correct option.
Assertion (A) : Deadlock avoidance ensures that THE system will never enter an unsafe state, from which a deadlock could occur.
Reason (R) : Deadlock avoidance schemes, like THE Banker's algorithm, require a priori information about THE maximum number of resources of each type that each process may need.
(A) Both A and R are true and R is THE correct explanation of A.
(B) Both A and R are true but R is not THE correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

191. A memory management unit (MMU) is responsible for:
(A) Translating logical addresses to physical addresses.
(B) Managing THE disk schedule.
(C) Performing context switches between processes.
(D) Caching frequently used data.

192. Which of THE following is an INCORRECT pair regarding disk structure?
(A) Track - Circular path on THE surface of a disk.
(B) Sector - Smallest unit of data that can be read or written.
(C) Cylinder - A set of tracks that are at THE same arm position on all platters.
(D) Seek Time - THE time for THE desired sector to rotate under THE head.

193. THE method of file allocation where all blocks of a file are stored in a non-sequential manner and linked together using pointers is called:
(A) Contiguous allocation
(B) Linked allocation
(C) Indexed allocation
(D) Grouped allocation

194. IN THE context of security, what is a 'trap door'?
(A) A program that appears legitimate but performs malicious actions.
(B) A hidden entry point in a program that bypasses normal security checks.
(C) An attack that floods a system with requests to make it unavailable.
(D) A self-replicating program that spreads across a network.

195. Fast User Switching in Windows is a feature that:
(A) Allows an administrator to switch to a different user's session without logging

Downloaded from www.pdf.tube

(A) Increasing THE number of processes in memory.
(B) Using a smaller page size.
(C) Providing a process with more memory frames (its working set).
(D) Using THE FIFO page replacement algorithm.

178. Match List-I with List-II regarding I/O concepts and select THE correct answer from THE options given below.
List I:(Concept)
P. I/O Hardware
Q. Application I/O Interface
R. Kernel I/O Subsystem
S. Directory Implementation
List II:(Example/Technique)
1. System calls (e.g., open, read, write)
2. Device drivers, DMA controller
3. Spooling, Caching
4. Hash Table
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-1, R-3, S-4
(D) P-3, Q-4, R-2, S-1

179. A capability list in an access control system is:
(A) A list for each object, specifying which domains can access it.
(B) A list for each domain, specifying THE objects and operations it can perform.
(C) A list of all users currently logged into THE system.
(D) A list of all available resources in THE system.

180. Virtualization allows a single physical machine to host multiple guest operating systems. This is managed by a piece of software known as a:
(A) Compiler
(B) Linker
(C) Hypervisor or Virtual Machine Monitor (VMM)
(D) Loader

181. IN THE Linux operating system, Interprocess Communication (IPC) can be achieved through which of THE following mechanisms?
(A) Pipes and FIFOs
(B) Message Queues and Semaphores
(C) Shared Memory
(D) All of THE above

182. A major design principle of modern Windows operating systems is its object-oriented nature, where system resources like files, processes, and devices are represented as:
(A) Basic data structures stored in THE kernel.
(B) Objects with associated attributes and services.
(C) Entries in THE persistent system registry.
(D) Set of dynamically linked system libraries.

NODIA

183. Read THE following statements and choose THE correct option.
Statement I : IN a distributed system, 'robustness' refers to THE ability to detect and recover from failures.
Statement II : A key design issue in distributed systems is 'transparency', which aims to hide THE fact that resources are physically distributed across a network.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

184. THE function of a linker is to:
(A) Place THE program in memory for THE purpose of execution.
(B) Relocate THE program to execute from a specific memory area.
(C) Combine various object modules into a load module.
(D) Translate a program from a high-level language to machine code.

185. Which of THE following is THE most appropriate definition of an operating system?
(A) A set of utility programs that manage hardware resources.
(B) A program that acts as an intermediary between THE user and hardware.
(C) A collection of compilers and interpreters.
(D) THE software that provides a graphical user interface.

devices.
(D) Cache data for I/O operations.

205. THE process of verifying THE identity of a user is called:
(A) Authorization
(B) Authentication
(C) Accounting
(D) Auditing

206. Read THE following statements and choose THE correct option.
Statement I : Para-virtualization requires THE guest operating system to be modified to run on THE hypervisor.
Statement II : Full virtualization allows an unmodified guest operating system to run in isolation.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

207. IN Linux, what is THE primary purpose of THE `scheduling` subsystem?
(A) To manage THE file system layout on THE disk.
(B) To allocate CPU time to runnable processes according to a scheduling policy.
(C) To handle I/O requests from user applications.
(D) To load kernel modules into memory.

208. Which is a correct pair regarding Windows OS components?
(A) HAL - Provides a graphical user interface.
(B) Executive - Manages core OS services like memory and I/O.
(C) Kernel - Sits on top of THE Executive and manages user profiles.
(D) File System - Manages interprocess communication.

209. One of THE major design issues for distributed systems is 'scalability'. This refers to THE system's ability to:
(A) Remain robust in THE face of failures.
(B) Handle more work by adding resources to THE system.
(C) Hide THE physical distribution of resources from users.
(D) Ensure secure communication between nodes.

210. What is a machine language?
(A) A high-level language that is easy for humans to read and write.
(B) A set of mnemonics that correspond to machine instructions.
(C) A set of binary-coded instructions that can be directly executed by a computer's CPU.
(D) A language used for communication between operating systems.

211. Which of THE following operating system structures is most likely to suffer from performance overhead due to frequent user/kernel mode switches for services?
(A) Monolithic Kernel
(B) Microkernel
(C) Layered System
(D) Modular Kernel

212. IN THE context of interprocess communication, what is a primary characteristic of message passing?
(A) It requires processes to share an address space.
(B) It is typically faster than shared memory.
(C) It is useful for communication across different computers.
(D) It is implemented solely by user-level libraries without OS support.

213. Arrange THE following storage devices in order of access speed, from fastest to slowest.
1. Main Memory (RAM)
2. Magnetic Disk
3. CPU Registers
4. Solid-State Disk (SSD)

Choose THE correct answer from THE options given below:
(A) 3, 1, 4, 2
(B) 1, 3, 4, 2
(C) 3, 4, 1, 2
(D) 4, 3, 1, 2

Click Here to Buy book on Amazon

them out.
(B) Quickly changes THE system's power plan.
(C) Switches between different network connections rapidly.
(D) Allows a user to run Linux applications directly on Windows.

196. THE system software that combines object files produced by a compiler into a single executable program is THE:
(A) Loader
(B) Linker
(C) Macro processor
(D) Debugger

197. Which of THE following is considered a service of THE operating system for ensuring efficient system operation?
(A) User interface
(B) Program execution
(C) Resource allocation
(D) File-system manipulation

198. Consider THE following statements about synchronization tools:
1. A semaphore is an integer variable accessed only through atomic `wait()` and `signal()` operations.
2. Peterson's solution is a hardware-based solution to THE critical section problem.
3. A mutex lock is essentially a binary semaphore.
4. Busy waiting can be avoided by using blocking semaphores where a process is put into a waiting queue.

Which of THE statements given above are correct?
(A) 1 and 2 only
(B) 1, 3, and 4 only
(C) 2 and 3 only
(D) 1, 2, 3 and 4

199. What is a primary issue with THE `fork()` system call in a multithreaded program?
(A) It always terminates THE parent process.
(B) THE question of whether it duplicates only THE calling thread or all threads.
(C) It can only be called by THE main thread.
(D) It is not supported by modern operating systems.

200. THE time it takes for THE dispatcher to stop one process and start another running is known as:
(A) Turnaround time
(B) Waiting time
(C) Dispatch latency
(D) Response time

201. Match List-I with List-II regarding deadlock concepts and select THE correct answer from THE options given below.
List I:(Concept)
P. Mutual Exclusion
Q. Hold and Wait
R. No Preemption
S. Circular Wait
List II:(Description)
1. A process holding a resource can request another, and must wait.
2. A chain of waiting processes P0, P1, ..., Pn where P0 waits for P1, etc., and Pn waits for P0.
3. A resource cannot be forcibly taken from a process.
4. At least one resource must be held in a non-sharable mode.
(A) P-1, Q-4, R-2, S-3
(B) P-4, Q-1, R-3, S-2
(C) P-4, Q-1, R-2, S-3
(D) P-1, Q-2, R-3, S-4

202. THE set of all logical addresses generated by a program is referred to as THE:
(A) Physical address space
(B) Logical address space
(C) Address binding
(D) Page table

203. Which of THE following is NOT a file access method?
(A) Sequential Access
(B) Direct Access
(C) Indexed Access
(D) Mounted Access

204. IN I/O systems, Direct Memory Access (DMA) is used to:
(A) Allow THE CPU to directly access I/O devices.
(B) Allow devices to transfer data to memory without CPU.
(C) Provide a memory-mapped interface for I/O

NODIA

machines.
(C) It is designed primarily for fault tolerance and high availability.
(D) All processing is done on a central server.

224. Which of THE following system software performs THE task of `relocation`?
(A) Compiler
(B) Assembler
(C) Linker or Loader
(D) Macro preprocessor

225. A `system call` provides an interface to THE services made available by:
(A) A compiler
(B) THE operating system
(C) THE system libraries
(D) An application program

226. Consider THE following activities related to process management:
1. Creating and deleting user and system processes.
2. Suspending and resuming processes.
3. Providing mechanisms for process synchronization.
4. Providing mechanisms for process communication.
5. Providing mechanisms for deadlock handling.

Which of these activities are THE responsibility of THE operating system's process management functions?
(A) 1 and 2 only
(B) 1, 2, and 3 only
(C) 1, 2, 3, and 4 only
(D) 1, 2, 3, 4, and 5

227. A `spin lock` is a type of mutex where a thread trying to acquire THE lock:
(A) Is put into a sleep state and placed in a waiting queue.
(B) Enters a loop, repeatedly checking if THE lock is available.
(C) Is immediately terminated.
(D) Is given a lower priority.

228. Which of THE following is a primary benefit of multithreading in a single process?
(A) Security, by isolating threads from each other.
(B) Simplicity, as it's easier to program than a single-threaded application.
(C) Responsiveness, by allowing a program to continue running even if part of it is blocked.
(D) Protection, by giving each thread a separate memory space.

229. Which of THE following scheduling algorithms is non-preemptive?
(A) Round Robin
(B) Shortest Remaining Time First (SRTF)
(C) First-Come, First-Served (FCFS)
(D) Preemptive Priority Scheduling

230. A resource-allocation graph is a directed graph used to model THE state of resource allocation. If this graph contains no cycles, then:
(A) THE system is in a deadlocked state.
(B) No process is in a waiting state.
(C) No deadlock exists in THE system.
(D) THE system is in an unsafe state.

231. THE concept of bringing a page into memory only when it is needed is called:
(A) Pre-paging
(B) Demand Paging
(C) Swapping
(D) Segmentation

232. Which RAID level combines mirroring and striping to provide both redundancy and performance?
(A) RAID 0
(B) RAID 1
(C) RAID 5
(D) RAID 10 (1+0)

233. THE file-system structure and implementation are made more portable across different file systems on a single OS by using:
(A) A Virtual File System (VFS) layer.
(B) A specific device driver for each file system type.
(C) Hard links and soft links.
(D) A single, universal file allocation method.

234. Read THE following statements and choose THE correct option.
Statement I : THE Access Matrix is a theoretical model of protection but is often impractical to

Downloaded from www.pdf.tube

157. A programmer writes code in C++, and then uses a g++ command to turn it into an executable file. This process involves, in order:
(A) Compilation, Linking, Loading
(B) Interpretation, Linking, Loading
(C) Linking, Compilation, Loading
(D) Compilation, Loading, Linking

158. THE 'dual-mode' operation (user mode and kernel mode) of an operating system is crucial for:
(A) Improving program execution speed.
(B) Protecting THE operating system from errant user programs.
(C) Allowing multiple users to log in simultaneously.
(D) Managing file storage on disk.

159. Consider THE following characteristics of a system:
1. Resources are allocated to processes before they begin execution.
2. A resource-allocation graph is used to check if granting a request leads to an unsafe state.
3. If a deadlock is found, a victim process is selected and terminated.
4. A total ordering is imposed on all resource types.
5. THE hold and wait condition is structurally disallowed.

Which of these are associated with Deadlock Prevention?
(A) 1, 4, and 5 only
(B) 2 and 3 only
(C) 1, 2, and 4 only
(D) 3 and 5 only

160. Which type of scheduling is responsible for controlling THE degree of multiprogramming?
(A) Short-term scheduler
(B) Medium-term scheduler
(C) Long-term scheduler
(D) I/O scheduler

161. What is THE primary purpose of a semaphore's V (signal) operation?
(A) To block a process until a resource is available.
(B) To decrement THE semaphore's value.
(C) To signal that a process is finished, potentially unblocking others.
(D) To initialize THE semaphore's value.

162. Read THE following assertion and reason carefully and choose THE correct option.
Assertion (A) : In THE One-to-One threading model, creating a user-level thread requires creating a corresponding kernel thread.
Reason (R) : This model can be a burden on THE system as creating a kernel thread is an expensive operation, potentially limiting THE total number of threads.
(A) Both A and R are true and R is THE correct explanation of A.
(B) Both A and R are true but R is not THE correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

163. A real-time operating system (RTOS) would be THE most appropriate choice for which of THE following systems?
(A) A desktop PC used for web browsing and word processing.
(B) A large-scale data processing batch system.
(C) An anti-lock braking system in a car.
(D) A file server for a small office network.

164. THE use of a page table in a memory management system is to:
(A) Map logical pages to physical frames.
(B) Store THE entire contents of a process.
(C) Keep track of free memory blocks.
(D) Implement a file allocation table.

165. THE strategy of always replacing THE page that was brought into memory first is called:
(A) LRU
(B) LFU
(C) FIFO
(D) Optimal

NODIA

166. Which of THE following disk scheduling algorithms may lead to starvation for some requests?
(A) FCFS
(B) SCAN
(C) C-SCAN
(D) SSTF (Shortest Seek Time First)

167. An access control list (ACL) is associated with:
(A) A process, specifying which objects it can

increases.
(C) THE anomaly where seek time increases with THE number of disk requests.
(D) THE anomaly where LRU performs worse than FIFO.

243. THE SSTF (Shortest Seek Time First) disk scheduling algorithm:
(A) Services requests in THE order they arrive in THE queue.
(B) Selects THE request with THE minimum seek time from THE current head position.
(C) Moves THE head from one end of THE disk to THE other, servicing requests along THE way.
(D) Is fair to all requests, regardless of their location on THE disk.

244. THE inode (index node) in Unix/Linux file systems is used to store:
(A) THE data of THE file.
(B) THE filename.
(C) THE attributes and disk block locations of THE file's data.
(D) THE directory path of THE file.

245. IN a distributed file system, location transparency means that:
(A) THE user knows THE exact physical location of a file.
(B) THE file name does not reveal THE file's storage location.
(C) THE file can be moved without its name changing.
(D) Multiple copies of THE same file are visible to THE user.

246. Which component is responsible for loading THE kernel of THE operating system from disk into memory?
(A) THE system call interface
(B) THE command interpreter (shell)
(C) THE bootstrap loader
(D) THE linker

247. THE 'progress' requirement for a solution to THE critical-section problem states that:
(A) A process waiting to enter its critical section will eventually be allowed to enter.
(B) If no process is in its critical section, selection of THE next cannot be postponed indefinitely.
(C) Only one process can be in its critical section at a time.
(D) A process must request permission before entering its critical section.

248. Which of THE following is an advantage of using threads over multiple processes?
(A) Threads are more secure because they are isolated from each other.
(B) Thread context switching is generally slower than process context switching.
(C) Threads share memory and resources, making communication and sharing easier.
(D) Each thread has its own independent address space.

249. IN indexed allocation for files, what is a major disadvantage if THE index block itself is too small?
(A) It leads to external fragmentation.
(B) It cannot support direct access.
(C) It limits THE maximum size of THE file.
(D) It makes sequential access very inefficient.

250. Read THE following statements and choose THE correct option.
Statement I : A Type 1 hypervisor is also known as a bare-metal hypervisor because it runs directly on THE host's hardware.
Statement II : A Type 2 hypervisor runs as an application on top of a conventional host operating system.
(A) Both Statement I and Statement II are correct.
(B) Both Statement I and Statement II are incorrect.
(C) Statement I is correct but Statement II is incorrect.
(D) Statement I is incorrect but Statement II is correct.

Direction : Please read the following source carefully.

System software serves as the foundational layer between computer hardware and application programs, facilitating the development and execution of software. This layer includes a spectrum of programming languages, from machine language, which is directly executed by the CPU, to assembly language, a low-level symbolic

Click Here to Buy book on Amazon

of Independent Disks) structure. The operating system provides a logical view of this storage through a file system, which abstracts the physical properties of devices. Files can be accessed via various access methods, such as sequential or direct access. The file system itself is organized using a directory and disk structure to manage the collection of files. To make a file system available, it must be mounted, a process that attaches the file system on a device to a specific location in the overall file system tree. Key implementation details include allocation methods—such as contiguous, linked, and indexed allocation—which determine how disk blocks are allocated to files. Efficiently managing available disk space is handled by free-space management techniques. The choice of these methods significantly impacts the system's overall efficiency and performance.

Based on the above source, attempt the following five questions.

286. What is the primary goal of disk scheduling algorithms?
(A) To organize the disk into tracks and sectors.
(B) To combine multiple disks into a single RAID structure.
(C) To manage the free space available on the disk.
(D) To minimize seek time and rotational latency for I/O requests.

287. The RAID structure is primarily used to:
(A) Define the access methods for files, such as sequential or direct.
(B) Mount a file system onto a specific location in the directory tree.
(C) Improve the performance and reliability of disk storage by combining multiple disks.
(D) Allocate disk blocks to files using contiguous or linked methods.

288. Which allocation method for files is mentioned in the passage?
(A) Sequential allocation
(B) Indexed allocation
(C) Layered allocation
(D) Mounted allocation

289. The process of making a file system on a device accessible within the system's file tree is known as:
(A) Free-space management
(B) Disk scheduling
(C) Directory structuring
(D) File system mounting

290. According to the passage, the file system provides a logical view of storage that abstracts the:
(A) Physical properties of the storage devices.
(B) Redundancy provided by the RAID structure.
(C) Access methods used by the applications.
(D) Directory structure of the entire system.

NODIA

Direction : Please read the following source carefully.

In a multi-user environment, an operating system must provide mechanisms for both protection and security. Protection refers to the mechanisms controlling the access of programs, processes, or users to the resources defined by a computer system. An abstract model for this is the Access Matrix, where rows represent subjects (e.g., users, processes) and columns represent objects (e.g., files, devices), with entries defining the access rights. This matrix is often implemented using access control lists or capability lists. A critical aspect of protection is the ability for revocation of access rights. Security, a broader concept, involves defending the system against internal and external attacks. These attacks manifest as various threats, which can be categorized into program threats (e.g., trojan horses, viruses) and system and network threats (e.g., worms, denial-of-service attacks). To counter these threats, cryptography has emerged as a crucial security tool, used for encrypting data to ensure confidentiality and integrity. Another fundamental security defense is user authentication, the process of verifying the identity of a user before granting them access to the system. Implementing a comprehensive security strategy involves deploying multiple layers of defense, including robust protection mechanisms, cryptographic protocols, and stringent user authentication policies to safeguard system resources and data.

Based on the above source, attempt the following five questions.

kernel. The overall design and implementation of an operating system can vary, from simple monolithic structures to more complex layered or microkernel approaches. The initial startup of a computer, known as the system boot, is a critical process managed by firmware. The boot process loads the operating system kernel into memory, which then initializes all system components, including device drivers and system services, ultimately preparing the system for user interaction and application execution.

Based on the above source, attempt the following five questions.

256. What is the primary role of the operating system as described in the passage?
(A) To execute application programs directly on the hardware without any mediation.
(B) To act as an intermediary that manages hardware and software resources for user applications.
(C) To provide a set of development tools for creating high-level language programs.
(D) To perform the initial startup sequence of the computer exclusively.

257. How do user applications typically request services like memory allocation or I/O device control from the operating system kernel?
(A) By directly accessing hardware memory addresses.
(B) By modifying the operating system's kernel code.
(C) By using compiler-specific intrinsic functions.
(D) By invoking system calls through a defined interface.

258. Which of the following is identified as the core of the operating system responsible for its essential functions?
(A) The kernel.
(B) The system boot loader.
(C) The set of application programs.
(D) The system call interface.

259. According to the passage, the system boot process is primarily responsible for:
(A) Managing interprocess communication and synchronization.
(B) Handling all system calls from user applications during runtime.
(C) Loading the operating system kernel into memory and initializing the system.
(D) Designing the structure and implementation of the operating system.

NODIA

260. The passage implies that a key reason for using system calls to access OS services is to:
(A) Ensure system integrity by providing controlled access to hardware.
(B) Improve the execution speed of user applications.
(C) Allow applications to choose between different operating system structures.
(D) Simplify the system boot process for faster startup.

Direction : Please read the following source carefully.

Process management is a central function of modern operating systems, involving the creation, scheduling, and termination of processes. For processes to collaborate, they require mechanisms for interprocess communication (IPC), which allow them to exchange data and information. A common model for IPC is the client-server system, where one process (the client) requests services from another process (the server). When multiple processes execute concurrently and share data, their operations must be synchronized to maintain data consistency. The challenge of ensuring that only one process can access a shared resource at a time is known as the critical-section problem. A segment of code where a process accesses shared resources is called a critical section. To solve this problem, various synchronization solutions have been developed. Peterson's Solution is a classic software-based solution for two processes that uses shared variables to ensure mutual exclusion. However, for more complex scenarios, hardware support or more robust OS-provided tools are necessary. Semaphores are a powerful synchronization tool that can be used to control access to a shared resource by multiple processes. A semaphore is essentially an integer variable that, apart from initialization, is accessed only through two standard atomic operations: wait() and signal(), which manage entry into and exit from the critical section.

Based on the above source, attempt the following five questions.

274. The passage suggests that the Shortest-Job-First (SJF) algorithm is optimal for minimizing which scheduling criterion?
(A) Average waiting time
(B) CPU utilization
(C) Throughput
(D) Response time

275. The practice of scheduling kernel-level threads rather than entire processes is referred to as:
(A) Real-time CPU scheduling
(B) Multiple-processor scheduling
(C) Thread scheduling
(D) Asymmetric multiprocessing

Direction : Please read the following source carefully.

In a multiprogramming environment, a deadlock is a critical state where a set of blocked processes each holds a resource and waits to acquire a resource held by another process in the set. A deadlock can be characterized by four necessary conditions that must hold simultaneously: mutual exclusion, where at least one resource is non-sharable; hold and wait, where a process holds a resource while waiting for another; no preemption, where a resource can only be released voluntarily by the process holding it; and circular wait, where a cycle of processes exists such that each process is waiting for a resource held by the next. There are several methods for handling deadlocks. Deadlock prevention aims to ensure that at least one of the four necessary conditions can never occur, for example, by requiring processes to request all resources at once to negate the 'hold and wait' condition. Deadlock avoidance requires the operating system to have advance information about the resources a process will need, allowing it to dynamically decide if granting a request could lead to an unsafe state. Alternatively, a system can allow deadlocks to occur and then employ a deadlock detection algorithm to identify them. Once a deadlock is detected, a recovery scheme, such as terminating one or more processes or preempting resources, must be used to break the deadlock.

Based on the above source, attempt the following five questions.

276. Which of the following is NOT one of the four necessary conditions for a deadlock to occur?
(A) Resource preemption
(B) Mutual exclusion
(C) Hold and wait
(D) Circular wait

277. The deadlock handling method that involves ensuring at least one of the necessary conditions for deadlock cannot hold is known as:
(A) Deadlock avoidance
(B) Deadlock detection
(C) Recovery from deadlock
(D) Deadlock prevention

278. According to the passage, what does the deadlock avoidance method require to function effectively?
(A) The ability to preempt resources from any process.
(B) Prior information about the future resource needs of processes.
(C) A policy that denies the mutual exclusion condition.
(D) A recovery scheme to terminate deadlocked processes.

279. What is the role of a deadlock detection algorithm?
(A) To prevent the 'circular wait' condition from ever occurring.
(B) To dynamically assess if granting a resource request is safe.
(C) To identify that a deadlock has already occurred in the system.
(D) To preempt resources from a process to give them to another.

280. If a system uses deadlock detection, what must follow after a deadlock is identified?
(A) A system reboot to clear all resource allocations.
(B) An algorithm to prevent future deadlocks from happening.
(C) A protocol to enforce the hold and wait condition.
(D) A recovery scheme to break the existing deadlock.

Direction : Please read the following source carefully.

Effective memory management is crucial for the performance of an operating system. Early

Downloaded from www.pdf.tube

268. Based on the passage, what is the primary purpose of multithreading models like many-to-one and one-to-one?
(A) To map user-level threads to kernel-level threads.
(B) To define the API for a thread library.
(C) To address threading issues like signal handling.
(D) To allocate a unique stack and register set for each thread.

269. Which of the following is identified in the passage as a challenge or threading issue in multithreaded programming?
(A) The mapping of user threads to kernel threads.
(B) The sharing of the code section among threads.
(C) The process of thread cancellation and signal handling.
(D) The use of a thread library to create threads.

270. What is a thread considered to be, according to the source?
(A) A heavyweight process with its own memory space.
(B) A tool exclusively managed by compilers.
(C) A model for relating user and kernel activities.
(D) The fundamental unit of CPU utilization.

Direction : Please read the following source carefully.

CPU scheduling is the basis of multiprogrammed operating systems, determining which process in the ready queue is to be allocated the CPU. The effectiveness of a scheduling policy is evaluated based on several scheduling criteria, including CPU utilization, throughput (processes completed per time unit), turnaround time (total time from submission to completion), waiting time (time spent in the ready queue), and response time (time from submission to the first response). Different CPU scheduling algorithms are designed to optimize different criteria. For instance, First-Come, First-Served (FCFS) is simple but can lead to high average waiting times, while Shortest-Job-First (SJF) is optimal for minimizing average waiting time but requires future knowledge. In systems that support multithreading, scheduling is often performed on kernel-level threads, a practice known as thread scheduling. The complexity increases with multiple-processor scheduling, where the scheduler must distribute the load among several CPUs. This can be approached through asymmetric multiprocessing, where one processor handles all scheduling, or symmetric multiprocessing (SMP), where each processor is self-scheduling. Furthermore, real-time CPU scheduling is essential for systems with timing constraints. It distinguishes between soft real-time systems, which prioritize critical processes without guaranteeing deadlines, and hard real-time systems, where tasks must be serviced by their deadline.

Based on the above source, attempt the following five questions.

NODIA

271. Which scheduling criterion measures the total time a process takes from its submission until its completion?
(A) Waiting time
(B) Response time
(C) Turnaround time
(D) Throughput

272. In the context of multiple-processor scheduling, what is the key characteristic of symmetric multiprocessing (SMP)?
(A) A single master processor handles all scheduling decisions for all other processors.
(B) Each processor is self-scheduling and takes processes from a common ready queue.
(C) Processes are permanently assigned to a specific processor for their entire lifetime.
(D) Scheduling is only performed on user-level threads, not kernel-level threads.

273. What is the primary distinction between hard and soft real-time systems as described in the passage?
(A) Hard real-time systems prioritize tasks, while soft real-time systems do not.
(B) Soft real-time systems use thread scheduling, while hard real-time systems use process scheduling.
(C) Soft real-time systems are for single-processor systems, while hard real-time is for multiprocessors.
(D) Hard real-time systems must meet task deadlines, whereas soft real-time systems only aim to.

representation, and high-level languages, which offer abstraction and portability. The translation from high-level languages to machine code is primarily handled by compilers and interpreters. A compiler translates the entire source program into an executable file before execution, while an interpreter translates and executes the program line by line. Once a program is compiled into object modules, they must be prepared for execution. The linking process combines these modules and resolves external references to create a single executable file. Subsequently, the loader places this executable into main memory. This loading may involve relocation, where the loader adjusts program addresses to reflect the actual memory locations assigned. To enhance programmer productivity, macros provide a mechanism for replacing a shorthand notation with a predefined sequence of statements during a pre-processing phase. For post-execution analysis and error correction, debuggers are indispensable tools that allow programmers to control program execution, inspect state, and identify faults.

Based on the above source, attempt the following five questions.

251. Which of the following best distinguishes a compiler from an interpreter based on their translation process?
(A) A compiler executes the program line by line, whereas an interpreter creates a complete executable file.
(B) A compiler translates the entire source program at once before execution, while an interpreter translates and executes code sequentially.
(C) An interpreter produces an object module for linking, whereas a compiler directly generates machine code for immediate execution.
(D) An interpreter is used for high-level languages, whereas a compiler is used exclusively for assembly languages.

252. According to the passage, what is the primary function of the linking process in software execution?
(A) To combine various object modules and resolve external references into a single executable.
(B) To load the executable file into main memory for execution.
(C) To translate high-level language code into machine-executable instructions.
(D) To adjust program addresses to match their physical location in memory.

253. The process of adjusting program addresses to reflect their final position in main memory is referred to as:
(A) Linking
(B) Loading
(C) Relocation
(D) Interpretation

254. Based on the description, a macro is best understood as a tool for:
(A) Identifying runtime errors after program execution.
(B) Translating assembly language into machine language.
(C) Combining multiple object files into a single executable program.
(D) Replacing a shorthand notation with a predefined block of code before compilation.

255. Which system software component is specifically designed to help programmers control program execution and inspect its state to identify faults?
(A) Compiler
(B) Loader
(C) Debugger
(D) Macro pre-processor

Direction : Please read the following source carefully.

The operating system (OS) is a fundamental piece of system software that manages all hardware and software resources of a computer. It provides a platform on which other programs can run, acting as an intermediary between the user and the computer hardware. The core of the OS, the kernel, performs essential functions through a set of operations and services, which include managing processes, allocating memory, handling file systems, and controlling I/O devices. To ensure system integrity and controlled access to hardware, user applications cannot perform these operations directly. Instead, they access OS services through a well-defined interface known as system calls. A system call is a programmatic way for a program to request a service from the

Click Here to Buy book on Amazon

systems used contiguous memory allocation, where each process is contained in a single section of memory. To improve utilization, swapping allows a process to be temporarily moved from main memory to a backing store and brought back for continued execution. More advanced techniques overcome the limitations of contiguous allocation. Paging divides physical memory into fixed-sized blocks called frames and logical memory into blocks of the same size called pages, allowing a process's physical address space to be non-contiguous. Segmentation, in contrast, supports the user's view of memory by dividing it into logical units of varying sizes, such as code, data, and stack. Virtual memory concepts further enhance efficiency. Demand paging is a widely used implementation of virtual memory where pages are loaded into memory only when they are needed, reducing I/O and memory requirements. However, this introduces the possibility of a page fault, requiring a page replacement algorithm to decide which page in memory to replace. If a process does not have enough frames, it can suffer from thrashing—a state of high paging activity where the system spends more time paging than executing. Memory-mapped files offer another sophisticated technique, allowing a portion of a file on disk to be mapped directly into a process's address space for I/O operations.

Based on the above source, attempt the following five questions.

281. What is the key difference between paging and segmentation?
(A) Paging uses fixed-sized blocks (pages/frames), whereas segmentation uses variable-sized logical units.
(B) Paging uses variable-sized blocks, while segmentation uses fixed-sized blocks.
(C) Paging divides memory into logical units, while segmentation divides it into physical frames.
(D) Paging is a form of virtual memory, whereas segmentation is not.

282. The technique of loading pages into memory only when they are required during execution is known as:
(A) Swapping
(B) Demand paging
(C) Contiguous memory allocation
(D) Segmentation

NODIA

283. What is thrashing in the context of memory management?
(A) A state where a process has too many frames allocated to it.
(B) A high level of paging activity where the system spends more time on I/O than on execution.
(C) The process of mapping a file directly into a process's address space.
(D) The procedure of replacing a page in memory when a page fault occurs.

284. The purpose of a page replacement algorithm is to:
(A) Select a page in memory to be replaced when a new page must be loaded.
(B) Decide which page to bring into memory from the backing store.
(C) Determine the appropriate frame size for physical memory.
(D) Prevent thrashing by allocating a sufficient number of frames.

285. How do memory-mapped files facilitate I/O operations?
(A) By swapping entire processes out to the file system.
(B) By allocating contiguous blocks of memory for each file.
(C) By mapping a disk file directly into a process's virtual address space.
(D) By dividing all files into fixed-size pages on the disk.

Direction : Please read the following source carefully.

Modern computer systems rely on a complex storage hierarchy, with mass-storage systems like magnetic disks forming a critical secondary storage layer. The physical disk structure is organized into platters, surfaces, tracks, and sectors. To optimize performance, the operating system must manage the hardware efficiently, which includes employing disk scheduling algorithms to service I/O requests in an order that minimizes seek time and rotational latency. For enhanced performance and reliability, multiple physical disks are often combined into a single logical unit using a RAID (Redundant Array

261. What is the primary purpose of interprocess communication (IPC) as described in the passage?
(A) To ensure that only one process can access a shared resource at a time.
(B) To schedule the execution of multiple processes on the CPU.
(C) To manage the creation and termination of processes.
(D) To allow cooperating processes to exchange data and information.

262. The segment of code in which a process accesses shared data is referred to as the:
(A) Semaphore
(B) Critical section
(C) Client-server system
(D) Peterson's Solution

263. According to the source, Peterson's Solution is characterized as a:
(A) Hardware-based solution for managing multiple client processes.
(B) General-purpose synchronization tool provided by the operating system.
(C) Software-based solution to the critical-section problem for two processes.
(D) Communication protocol used in client-server systems.

264. How does a semaphore function as a synchronization tool?
(A) By allowing any number of processes into a critical section simultaneously.
(B) By creating a client-server relationship between communicating processes.
(C) By dividing a program into independent processes for parallel execution.
(D) By using a shared integer variable and atomic operations to control access to a resource.

265. The critical-section problem arises due to:
(A) The concurrent access to shared resources by multiple processes.
(B) The need for communication between a client and a server.
(C) The requirement for processes to be scheduled by the operating system.
(D) The termination of a process before its execution is complete.

Direction : Please read the following source carefully.

The advent of multicore processors has made parallel execution a mainstream programming paradigm, with threads being the fundamental unit of CPU utilization. A thread, or lightweight process, comprises a thread ID, a program counter, a register set, and a stack, but shares its code section, data section, and other operating-system resources with other threads belonging to the same process. To manage the relationship between user-level threads (created by programmers) and kernel-level threads (managed by the OS), different multithreading models are used. These include the many-to-one model, one-to-one model, and many-to-many model. Programmers typically create and manage threads using a thread library, which provides a specific API. However, a modern approach known as implicit threading transfers the responsibility of creating and managing threads from the application developer to the compilers and run-time libraries. This allows developers to focus on application logic while the system handles the complexities of parallelism. Despite the benefits, multithreading introduces several challenges, referred to as threading issues. These issues include problems related to thread cancellation, signal handling in a multithreaded environment, and maintaining thread-specific data, all of which require careful design and management for correct and efficient parallel program execution.

Based on the above source, attempt the following five questions.

266. What key resources does a thread share with other threads within the same process?
(A) Its own unique program counter and register set.
(B) The process's code section and data section.
(C) A separate stack and thread ID.
(D) The management of the thread library.

267. Which concept involves shifting the creation and management of threads from the programmer to compilers and run-time libraries?
(A) Multicore programming
(B) Implicit threading
(C) Thread libraries
(D) Multithreading models

Downloaded from www.pdf.tube

291. What is the primary difference between protection and security as described in the source?
(A) Protection deals with external threats, while security deals with internal threats.
(B) Protection uses cryptography, while security uses access control lists.
(C) Protection is a mechanism to control resource access, while security is a broader defense against all attacks.
(D) Protection is about user authentication, while security is about program threats.

292. The Access Matrix is described as an abstract model where entries specify:
(A) The access rights of a subject over an object.
(B) The threats posed by a subject to an object.
(C) The cryptographic keys used to access an object.
(D) The location of an object within the file system.

293. Which of the following is categorized as a system and network threat?
(A) A trojan horse
(B) A virus
(C) An access matrix vulnerability
(D) A denial-of-service attack

294. The passage identifies cryptography as a security tool primarily used for:
(A) Implementing the Access Matrix and revoking rights.
(B) Encrypting data to ensure confidentiality and integrity.
(C) Authenticating user identities through passwords.
(D) Controlling access between subjects and objects.

295. What is the main purpose of user authentication in the context of system security?
(A) To encrypt all data transmitted over the network.
(B) To define the access rights a user has on different files.
(C) To verify a user's identity before allowing access to the system.
(D) To detect and remove program threats like viruses.

NODIA

Direction : Please read the following source carefully.

The design principles of major operating systems like Linux and Windows illustrate different architectural philosophies. Linux is known for its monolithic kernel design, which integrates most services into the main kernel space. However, it maintains flexibility through kernel modules, which can be dynamically loaded and unloaded at runtime, allowing for features like new file systems or device drivers to be added without a reboot. Its process management, scheduling, and memory management adhere to established UNIX principles. In contrast, the Windows operating system is based on a design that is more akin to a microkernel, with many OS services running as system components in user mode, though it is not a pure microkernel. Key components include its support for terminal services and fast user switching, and its advanced NT File System (NTFS). Moving beyond single-machine systems, distributed systems consist of multiple autonomous computers that communicate through a network. These are often built on network-based operating systems that manage resources across the network. The communication structure and its underlying protocols are vital for coordination. Significant design issues for distributed systems include achieving robustness, which is the ability to tolerate faults in individual nodes, as well as ensuring transparency and scalability, which are far more complex challenges than in a centralized OS.

Based on the above source, attempt the following five questions.

296. What feature of the Linux kernel provides flexibility despite its monolithic design?
(A) Its adherence to established UNIX principles.
(B) The use of terminal services for remote access.
(C) Its process management and scheduling algorithms.
(D) The ability to dynamically load and unload kernel modules.

297. How does the architectural design of Windows differ from that of Linux, according to the passage?
(A) Windows runs many OS services as system

components, similar to a microkernel approach.
(B) Windows uses a pure monolithic kernel, while Linux uses a microkernel.
(C) Linux supports fast user switching, whereas Windows does not.
(D) Windows uses dynamically loadable modules, whereas Linux does not.

298. The ability of a distributed system to continue functioning despite failures in some of its nodes is referred to as:
(A) Scalability
(B) Robustness
(C) Transparency
(D) Communication structure

299. Which of the following is identified as a key component or feature specific to the Windows operating system in the passage?
(A) A monolithic kernel design.
(B) The NT File System (NTFS).
(C) Kernel modules for dynamic loading.
(D) Adherence to UNIX design principles.

300. The passage implies that a major challenge in designing distributed systems compared to centralized ones is:
(A) Managing communication structures and achieving fault tolerance.
(B) Implementing process and memory management.
(C) Designing a suitable file system like NTFS.
(D) Creating a user interface for fast user switching.

Click Here to Buy book on Amazon

ANSWERS

1. (B) To transfer an executable file from secondary storage into main memory for execution.
Loaders are system programs responsible for preparing executable programs for execution by placing them in main memory. This involves taking the binary executable file from secondary storage, such as a hard drive, and mapping it into the physical address space. Proper loading ensures that the CPU can access instructions and data during runtime operations.

2. (A) The standard programmatic method for a user process to request a kernel-level service.
A system call acts as the fundamental interface between a running application and the operating system kernel. It provides a controlled mechanism for user-mode processes to request privileged operations, such as file management, process creation, or network access. This transition from user mode to kernel mode ensures system security and provides stable resource management.

3. (D) 1, 2, 3, and 4
System software encompasses various tools essential for program execution and development. Interpreters process code line-by-line for immediate execution, while compilers transform entire source files into object code. Relocation adjusts internal memory addresses to allow software to run at different locations. Macros simplify coding by allowing short identifiers to represent longer sequences of computer instructions.

4. (C) Statement I is correct but Statement II is incorrect.
Peterson's solution is a classic software-based algorithm designed to solve the critical-section problem for two processes. It ensures mutual exclusion, progress, and bounded waiting through shared variables. However, the solution is restricted to exactly two processes and relies on busy waiting, where a process continuously checks a condition, consuming processor cycles instead of yielding.

5. (C) Each thread has its own independent process control block (PCB).
Threads are lightweight units of execution within a process that share the same address space and resources. While processes have their own Process Control Blocks for management, threads typically have smaller, specialized structures called Thread Control Blocks. Context switching between threads is significantly more efficient than switching between processes due to the shared virtual memory.

6. (A) P-4, Q-1, R-2, S-3
CPU scheduling involves choosing the next process for execution based on specific criteria. First-Come First-Served can lead to the convoy effect, while Shortest Job First is optimal for minimizing average waiting time. Round Robin is tailored for time-sharing environments, and Priority Scheduling allows for task importance but may result in the starvation of processes.

7. (D) Deadlock
A deadlock represents a specific stall in a multiprogramming system where a set of processes is permanently blocked. This occurs when every process in the set holds at least one resource and is simultaneously waiting for another resource held by another process. This circular dependency prevents any further progress without intervention from the operating system.

8. (A) Thrashing
Thrashing occurs in virtual memory systems when the operating system spends a disproportionate amount of time swapping pages between main memory and secondary storage. This happens when the total demand for memory frames exceeds the available physical memory, causing a drastic drop in CPU utilization because the processor is idle while waiting for disk operations.

9. (C) A is true but R is false.
RAID Level 1 provides data redundancy and high reliability through a process known as mirroring. In this configuration, identical data is written to

Downloaded from www.pdf.tube

two or more separate disks simultaneously. While this improves read performance by allowing simultaneous access, it does not use data striping, which is a characteristic of RAID Level 0 for parallel performance.

10. (A) Direct access
Direct access, often called random access, is a file access method where records can be read or written in any order. This is achieved by specifying the exact block number or address of the data on the storage device. It is highly efficient for databases and applications requiring frequent, non-sequential data retrieval from various locations.

11. (D) Application software compilation
Operating systems provide essential services like program execution, input and output operations, and file system management. They also handle communications and error detection. However, the compilation of application software is performed by a compiler, which is a separate development tool. Compilation is not considered a core runtime service provided directly by the operating system kernel.

12. (B) 1, 2, 3, and 5 only
Security defenses in an operating system utilize various mechanisms to protect resources and verify identities. Access matrices define permissions, while cryptography secures data transmission. User authentication ensures that individuals are who they claim to be. Revocation of rights is a management process to remove access. Program threats, however, represent the actual dangers or attacks themselves.

13. (C) Type 2 (hosted) hypervisor
A Type 2 hypervisor is a virtualization layer that runs as an application on top of an existing host operating system. It relies on the host OS to manage physical hardware resources like memory and processor time. This structure is common in desktop environments where users need to run multiple operating systems alongside their primary workstation.

14. (D) To load hardware drivers and system functions into the kernel dynamically as needed.
Kernel modules in Linux are specifically designed to extend the functionality of the kernel without requiring a full system reboot. They allow the operating system to load and unload code, such as hardware drivers or filesystem support, on demand. This modular approach keeps the core kernel small while maintaining flexibility for diverse hardware and software requirements.

15. (B) 4, 2, 1, 3
The boot process begins with the Power-On Self-Test to verify hardware integrity. Following this, the BIOS or UEFI initiates the bootstrap loader, which locates and loads the operating system kernel into the main memory. Once the kernel is active, it initializes system services and starts background daemons to prepare the system for user interaction.

16. (B) Hardware Abstraction Layer (HAL)
The Hardware Abstraction Layer serves as an intermediary between the Windows operating system and the physical hardware. It hides hardware-specific details from the rest of the kernel, allowing higher-level software to interact with components through standardized interfaces. This design enhances portability, as the same core OS can run on diverse hardware platforms with different motherboard architectures.

17. (B) Ensuring robustness and fault tolerance across a network
Distributed systems are composed of multiple independent computers that appear as a single system. A primary design challenge is maintaining robustness and fault tolerance, ensuring the system remains operational even if specific nodes or network links fail. This requires complex coordination and synchronization mechanisms that are not as critical in traditional centralized or uniprocessor operating systems.

18. (D) 1, 2, 3, and 4
Efficient process management is vital for modern operating systems. The ready queue stores processes prepared for execution, while the long-term scheduler manages the degree of multiprogramming. Context switching enables multitasking by saving and restoring process states. Additionally, interprocess communication provides essential mechanisms for data exchange between active tasks using shared memory or message-passing protocols.

NODIA

19. (C) Multiple user-level threads are mapped to a single kernel thread.
In the Many-to-One multithreading model, a thread library manages multiple user-level threads which are then multiplexed onto one kernel-level thread. While this approach is efficient for thread management in user space, it lacks true parallelism on multiprocessor systems. Furthermore, if one thread performs a blocking system call, the entire process and all its threads are halted.

20. (D) Shortest Job First (SJF)
The Shortest Job First scheduling algorithm prioritizes processes with the smallest CPU burst times to minimize average waiting time. While effective for overall efficiency, it can lead to the starvation of longer processes if short tasks continuously arrive in the ready queue. Longer jobs may never receive CPU time as they are consistently bypassed by newer, shorter requests.

21. (B) P-3, Q-4, R-1, S-2
Deadlock handling involves several distinct strategies. Prevention ensures that at least one necessary condition for deadlock cannot occur. Avoidance uses algorithms like the Banker's algorithm to maintain safe states. Detection identifies deadlocks through cycles in wait-for graphs. Finally, recovery involves resolving the deadlock by terminating processes or preempting resources to allow other tasks to proceed.

22. (B) To store recently used page table entries for faster address translation.
The Translation Look-aside Buffer is a specialized, high-speed hardware cache used by the memory management unit. It stores recent mappings from virtual page numbers to physical frame numbers. By checking the TLB before the main page table in memory, the system significantly reduces the time required for address translation, improving the overall execution speed of processes.

23. (A) Both Statement I and Statement II are correct.
Memory management techniques address different types of inefficiencies. Contiguous allocation requires a single continuous block of memory, which often leads to external fragmentation where free space is scattered and unusable. Paging solves this by dividing processes into fixed-size pages that can be placed in non-contiguous frames, effectively eliminating external fragmentation while introducing minor internal fragmentation.

24. (A) It services requests in one direction and returns to start without servicing the return trip.
C-SCAN, or Circular SCAN, is a disk scheduling algorithm designed to provide a more uniform wait time than the standard SCAN algorithm. It services requests in one direction until it reaches the end of the disk. Instead of servicing requests on the way back, the head immediately returns to the beginning of the disk to start again.

25. (C) Access Matrix
An access matrix is a formal security model used to define the permissions and access rights of various subjects over specific objects. The rows represent domains or subjects, while the columns represent objects like files or devices. Each entry in the matrix specifies the set of operations, such as read or write, that a domain can perform.

26. (A) Kernel Modules
Interprocess communication in Linux is facilitated through several standard mechanisms including pipes for data streams, message queues for discrete data packets, and semaphores for synchronization. Shared memory is another common method. Kernel modules, however, are used to extend the functionality of the operating system itself and are not primarily used as a mechanism for communication between user-level processes.

27. (D) Fast User Switching
Fast User Switching is a convenient feature in Windows that allows multiple users to remain logged into their accounts simultaneously. This mechanism enables one person to use the computer while another's applications and files continue running in the background. It improves productivity by eliminating the need for users to close all programs and log out before another starts.

28. (B) Loader - Modifies a program to run from a new memory location
Compilers and interpreters translate source code, while linkers combine object modules into an executable. The primary job of a loader is to place the executable program into main memory for

Click Here to Buy book on Amazon

execution. Relocation, which involves modifying a program to run from a new memory location, is often handled by the linker or a specialized part of the loader.

29. (C) The OS provides services for both users and the system, like resource allocation and accounting.
Operating systems offer a wide variety of services tailored to both user convenience and efficient system administration. These include managing user interfaces, executing programs, and facilitating communications. From a system perspective, the operating system ensures optimal resource allocation among multiple processes, maintains system accounting, and implements robust security and protection measures to prevent unauthorized access and resource interference.

30. (A) 1, 2, and 4 only
Process synchronization involves managing access to shared resources to prevent data inconsistency. A race condition occurs when concurrent access results in unpredictable outcomes. The critical-section problem aims to coordinate these accesses. While semaphores are a key synchronization tool, they are implemented in software and hardware. Peterson's Solution is an elegant software method, but it only applies to two processes.

31. (A) Creating and managing threads at the user level.
Thread libraries like Pthreads provide a standardized Application Programming Interface for programmers to create and manage threads within their software. These libraries handle the complexities of thread lifecycle, synchronization, and communication. Depending on the underlying operating system model, these user-level threads may be mapped to kernel threads, allowing developers to implement concurrency without managing hardware-level details.

32. (B) Turnaround time
Turnaround time is a crucial performance metric in CPU scheduling that measures the total elapsed time from a process's initial submission to its final completion. This includes the time spent waiting to enter memory, waiting in the ready queue, executing on the processor, and performing input/output operations. It provides a comprehensive view of how long a process resides.

33. (A) Both A and R are true and R is the correct explanation of A.
Deadlock prevention strategies involve negating at least one of the four necessary conditions for a deadlock. To address the hold and wait condition, the system can mandate that a process must request and receive all its required resources before it starts executing. This ensures that a process never holds resources while simultaneously waiting for additional ones, thereby preventing cycles.

34. (A) Swapping
Swapping is a memory management technique used to increase the degree of multiprogramming. When physical memory becomes full, the operating system moves entire processes or large portions of them from main memory to a dedicated backing store on the disk. The process is later brought back into memory when it is scheduled to continue its execution on the processor.

35. (C) RAID 5
RAID 5 utilizes block-level striping with distributed parity to provide a balance of storage efficiency, performance, and data protection. Parity information is spread across all disks in the array rather than stored on a dedicated disk. This configuration allows the system to recover data if any single disk fails, while maintaining high read performance through parallel access.

36. (D) Tree-structured directory structure
A tree-structured directory is the most common organization in modern file systems. It consists of a root directory with multiple levels of subdirectories and files. This hierarchical arrangement allows users to organize data logically and prevents filename conflicts between different directories. It provides a flexible and intuitive way to manage large numbers of files across various user accounts.

37. (C) Trojan horse
A Trojan horse is a deceptive form of malware that presents itself as a useful or legitimate program. Once executed, it performs unauthorized or

NODIA

malicious actions in the background, such as stealing data or providing back-door access to the system. Unlike worms or viruses, a Trojan horse does not typically self-replicate but relies on social engineering for distribution.

38. (B) It runs directly on the hardware to control it and manage guest systems.
Type 1 hypervisors, also known as bare-metal hypervisors, are installed directly on the physical hardware of a machine. They operate with high privilege to manage hardware resources and coordinate multiple guest operating systems. This architecture is common in enterprise server environments because it offers superior performance and security compared to hosted hypervisors that run on a host OS.

39. (D) It offers high performance for kernel communication while maintaining flexibility.
The Linux kernel uses a monolithic architecture where all core services run in a single address space for maximum performance. To overcome the rigidity of traditional monolithic designs, it supports loadable kernel modules. This allows the system to remain flexible by loading only the necessary drivers and functions at runtime, thus optimizing memory usage while maintaining high-speed internal communication.

40. (A) They are designed to provide a high degree of transparency and fault tolerance.
Network operating systems allow multiple autonomous computers to connect and share resources. Unlike distributed operating systems, which aim for high transparency so that the network appears as a single machine, network OS users are explicitly aware of the different machines. Each machine maintains its own local operating system, and accessing remote resources usually requires specific, non-transparent user commands.

41. (C) Assembly Language
Assembly language is a low-level programming language that provides a human-readable representation of machine instructions. It uses short mnemonics to represent binary operations and symbolic names for memory addresses. While it is much easier to read than binary machine code, it still requires an assembler to translate the mnemonics into the final executable machine language for the CPU.

42. (A) It creates a new process that is a duplicate of the calling process.
The `fork()` system call is a fundamental operation in UNIX-like systems used to create a new process. The new process, called the child, is an exact copy of the parent process, including its memory contents and file descriptors. Both processes continue execution from the instruction immediately following the call, usually distinguished by different return values from the function.

43. (B) 2, 3, and 5 only
Starvation occurs when a process is indefinitely delayed because higher-priority or shorter tasks are always favored by the scheduler. Preemptive and non-preemptive Shortest Job First can cause starvation for long processes. Priority scheduling also suffers from this if high-priority jobs continually enter the system. First-Come First-Served and Round Robin are generally considered fair and do not cause starvation.

44. (A) Both Statement I and Statement II are correct.
Semaphores are synchronization variables used to manage access to shared resources through two atomic operations. The `wait()` operation, or P, decrements the semaphore value, signaling that a process is requesting a resource. The `signal()` operation, or V, increments the value, indicating a resource has been released. These operations ensure mutual exclusion and coordinate process activities in concurrent systems.

45. (C) A blocking system call by one thread does not block the entire process.
In the One-to-One multithreading model, every user-level thread is mapped to its own kernel thread. This provides a significant advantage over the Many-to-One model because if one thread performs a blocking system call, the operating system can schedule other threads from the same process. This allows for higher concurrency, especially on multi-core systems where multiple threads run.

46. (A) 3, 2, 1, 4
Handling an I/O request starts when a user process makes a system call. The kernel receives

Downloaded from www.pdf.tube

ends up with small, unusable gaps between segments. This phenomenon, known as external fragmentation, can prevent the system from loading new segments even if the total free space is sufficient.

84. (A) To provide a standardized interface between device drivers and the rest of the kernel.
The kernel I/O subsystem acts as an abstraction layer within the operating system. It provides a uniform interface that allows different parts of the kernel to communicate with diverse device drivers. By standardizing these interactions, the subsystem simplifies the management of I/O devices and enables features like buffering, caching, and spooling, which improve the overall efficiency of data transfer operations.

85. (D) Access-matrix-based (e.g., permission tables)
User authentication is the process of verifying a person's identity before allowing system access. Common methods include something the user knows, like a password, something they have, such as a smart card, or something they are, like a fingerprint. An access matrix, however, is a model for authorization, which defines what an already authenticated user is permitted to do.

86. (A) It reduces the physical hardware requirements for running multiple operating systems.
Virtualization provides a way to run multiple, isolated operating systems on a single physical server. By abstracting the hardware, it allows for better resource utilization and reduces the need for multiple physical machines. This leads to significant savings in power, space, and hardware costs, making it a fundamental technology for modern data centers and cloud computing environments.

87. (B) The set of services exposed by the Executive
In the Windows architecture, the Executive provides a broad set of services that applications can call to perform system-level tasks. These services are the Windows equivalent of system calls in UNIX or Linux. They provide a controlled interface for user-mode applications to interact with the kernel for operations like process management, memory allocation, security checks, and input/output management.

88. (B) A system that allows clients to access files stored across several network servers.
A Distributed File System is a storage architecture that allows users to access and manipulate files located on remote servers as if they were stored locally. The primary goal is to provide a unified view of data distributed across multiple machines in a network. This improves data availability, allows for easier sharing among users, and can provide increased storage capacity.

89. (C) Least Recently Used (LRU)
The Least Recently Used page-replacement algorithm selects the page that has not been accessed for the longest time to be replaced. This policy is based on the principle of temporal locality, assuming that pages used recently are likely to be used again soon. While more complex to implement than FIFO, it generally provides a much lower page-fault rate in practice.

90. (D) A necessary but not sufficient condition for a deadlock if there are multiple instances of a resource type.
In a resource-allocation graph, a cycle indicates a circular dependency. If each resource type has only one instance, a cycle is a definite indicator of a deadlock. However, if resource types have multiple instances, a cycle might exist without a deadlock, as a process outside the cycle could release a resource that resolves the dependency. Thus, it is necessary but insufficient.

91. (C) P-4, Q-3, R-1, S-2
Operating system concepts are linked to specific mechanisms. The system boot process starts with a bootstrap loader. Interprocess communication often utilizes message passing for data exchange. Process synchronization frequently employs semaphores to coordinate access to shared resources. The critical-section problem, a core synchronization challenge, can be addressed for two processes using the classic software-based Peterson's Solution.

92. (B) Resolving external symbols and combining multiple object files.
A linker is a system program that takes one or more object files produced by a compiler and combines them into a single executable file. Its primary task is symbol resolution, which involves connecting references to functions and variables

NODIA

56. (B) The system can allocate resources in some order and still avoid a deadlock.
In the Banker's algorithm, a state is safe if there exists a specific sequence of process executions that allows every process to eventually satisfy its maximum resource requirements and complete. This safe sequence ensures that even in the worst-case scenario where everyone requests their maximum resources at once, the system can still prevent a deadlock by ordering allocations.

57. (A) Both A and R are true and R is the correct explanation of A.
Paging is a memory management scheme that divides physical memory into fixed-size frames. When a process is loaded, its last page often does not perfectly match the frame size, leaving unused space at the end of the allocated block. This wasted space within a frame is known as internal fragmentation. While paging eliminates external fragmentation, this internal waste is unavoidable.

58. (B) To make a file system on a storage device accessible to the OS at a specific location in the directory tree.
File system mounting is the process of attaching an additional file system, such as one on a USB drive or a network share, to the primary directory structure. The operating system assigns a specific directory, called a mount point, through which the contents of the new file system become accessible. This allows the system to present a unified view of all data.

59. (D) All of the above.
Revocation involves removing access rights from a user or process. Immediate revocation can be achieved by scanning access lists. Periodic revocation automatically removes rights after a duration. Indirect methods use pointers to a central capability that can be invalidated. All these strategies provide different trade-offs between implementation complexity and the speed at which the rights are actually removed from the system.

60. (C) Executive
The Windows Executive consists of a set of kernel-mode components that provide core services such as memory management, process and thread management, and security. It sits above the Hardware Abstraction Layer and the kernel. The Executive manages the complex high-level tasks of the operating system, while the kernel itself handles lower-level scheduling and synchronization for the system's processors.

61. (C) High-Level Language
High-level languages, such as Python or Java, are designed for human readability and abstract away hardware details. Because CPUs only understand machine language consisting of binary code, these high-level programs must be translated. This is done by a compiler, which translates the whole program, or an interpreter, which translates and executes the code line by line during the program's runtime.

62. (D) Relocation
Relocation is the process performed by a linker or loader to map program addresses to actual physical memory locations. During compilation, addresses are often relative. Relocation adjusts these symbolic or relative references so they point to the correct spots in RAM where the program is actually loaded. This allows multiple programs to be loaded into different memory areas without conflicts.

63. (D) A monolithic kernel is always slower than a microkernel due to its large size.
Operating system architecture impacts performance and complexity. A monolithic kernel runs all services in kernel space, providing high-speed internal communication. In contrast, a microkernel moves many services to user space, which can introduce overhead due to frequent context switching. Therefore, monolithic kernels are often faster because they avoid the communication delays inherent in the modular, user-level approach of microkernels.

64. (B) Remote Procedure Calls (RPC)
Remote Procedure Call is a powerful communication mechanism that allows a process on one computer to execute a function on another computer. The complexity of network communication is hidden, so the programmer can write the code as if the procedure were local. This abstraction is fundamental for building distributed applications where different components reside on separate servers across a network.

Click Here to Buy book on Amazon

this request and uses the device driver to load specific commands into the device controller's registers. The controller then manages the hardware to perform the operation. Once the physical task is finished, the controller generates an interrupt to notify the CPU of completion.

47. (A) Hold and Wait
Hold and wait is one of the four necessary conditions for a deadlock. It describes a situation where a process has already acquired at least one resource and is currently keeping it while simultaneously waiting to obtain more resources held by other processes. This behavior creates dependencies between tasks that can eventually lead to a circular wait and a system stall.

48. (B) The process tries to access a page that is not currently in main memory.
A page fault is a type of interrupt raised by the hardware when a program attempts to access a block of memory that is part of its virtual address space but is not currently loaded in physical RAM. When this happens, the operating system must locate the required page on the disk and bring it into a free memory frame.

49. (B) Indexed allocation
Free-space management tracks the available blocks on a storage device. Common methods include bit vectors, which use a sequence of bits to represent block status, and linked lists that connect free blocks. Grouping is another technique that stores addresses of free blocks. Indexed allocation, however, is a file allocation method used to link a file's data blocks, not manage free space.

50. (A) Symmetric cryptography
Symmetric cryptography, also known as private-key encryption, uses a single shared key for both the encryption of plaintext and the decryption of ciphertext. Both the sender and the receiver must possess this secret key. While it is computationally faster than asymmetric methods, it requires a secure way to distribute and manage the secret key among all authorized parties involved.

51. (B) P-3, Q-4, R-1, S-2
Operating system components are linked to specific techniques. The kernel I/O subsystem manages efficiency through buffering and spooling. Directory implementation uses data structures like hash tables. The system boot process relies on a bootstrap loader to start the kernel. For client-server communication, sockets or Remote Procedure Calls are utilized to exchange data between different processes across a network infrastructure.

52. (C) To help a programmer find errors (bugs) in a program.
A debugger is a specialized software tool used during the development process to test and troubleshoot programs. It allows developers to execute code line by line, inspect variable values, and monitor the program's state at various points. This level of control helps identify the root causes of logical errors or crashes, making the software development process more efficient and reliable.

53. (C) Microkernel approach
The microkernel approach structures the operating system by minimizing the code running in the privileged kernel mode. Only essential services like interprocess communication and basic memory management remain in the kernel. All other services, such as device drivers and file systems, run as user-level processes. This modularity improves system reliability and security, as a failure in one service doesn't crash the kernel.

54. (B) 1, 2, and 3 only
Threads are efficient for multitasking because they are lightweight and share the same address space as their parent process. This sharing makes communication between threads of the same process faster than interprocess communication. Since they share memory, any modification to a global variable by one thread is immediately visible to all other threads in that same process, facilitating data exchange.

55. (C) Hard real-time task
A hard real-time task is a process that must be completed within a strict, predefined time limit. Failure to meet this deadline can lead to catastrophic system failure or unacceptable outcomes. These tasks are critical in systems like flight controllers or medical equipment, where timing precision is vital for safety. The scheduler must guarantee that these tasks always meet their deadlines.

NODIA

computers function like a single uniprocessor for the user. In contrast, network operating systems are less transparent, as users are typically aware of the individual machines on the network. Each machine in a network OS runs its own local operating system and manages its own resources and user accounts independently.

75. (B) To expand macro definitions by replacing macro calls with the macro's body.
A macro preprocessor is a tool that processes source code before compilation. It looks for macro calls, which are shorthand notations for a sequence of instructions. The preprocessor replaces each call with the full code body defined for that macro. This technique helps programmers write more concise and readable code by automating the repetition of frequently used instruction patterns.

76. (A) 4, 2, 3, 1
Memory management has evolved significantly over time. Early systems used contiguous allocation with fixed partitions. As multiprogramming grew, swapping was introduced to move processes between memory and disk. Later, segmentation allowed programs to be divided into logical sections. Eventually, paging became the standard, allowing for non-contiguous allocation and sophisticated virtual memory management to handle increasingly complex software and larger process sizes.

77. (B) An interrupt or scheduler dispatch
A process moves from the running state to the ready state when it is interrupted by the system or preempted by the scheduler. This typically happens in time-sharing systems when a process's assigned time quantum expires. The process remains ready to execute but must wait for the dispatcher to allocate the CPU again, allowing other tasks to have their turn.

78. (C) Mutex lock
A mutex lock, short for mutual exclusion, is a simple synchronization tool used to protect critical sections of code. It functions as a binary semaphore with a value of either zero or one. A process must acquire the lock before entering a critical section and release it afterward. This ensures that only one process can access a shared resource at any time.

79. (D) 1, 2, 3, and 4
CPU scheduling performance is measured by several factors. Turnaround time encompasses the entire life of a process, while response time measures the speed of first interaction. The convoy effect is a specific bottleneck in First-Come First-Served scheduling. Preemptive scheduling is a technique that allows the operating system to interrupt a running process to ensure that higher-priority tasks can be executed.

80. (B) A situation where outcome depends on the order of execution of multiple threads.
A race condition is a critical synchronization issue that occurs when multiple processes or threads concurrently access and manipulate shared data. The final state of the data depends on the exact order in which the operations are performed. Without proper coordination, such as using semaphores or mutexes, race conditions can lead to data corruption, logical errors, and unpredictable program behavior.

81. (B) Implicit threading
Implicit threading is a technique where the management of threads is handled by the compiler or a runtime library rather than the application programmer. Thread pools are a common example, where a set of threads is created at startup and assigned tasks as they arrive. This approach simplifies development and can improve performance by reducing the overhead of frequent thread creation.

82. (A) Both A and R are true and R is the correct explanation of A.
The Banker's algorithm is a deadlock avoidance strategy that requires each process to declare the maximum amount of each resource type it might ever need. This information is essential because the algorithm must simulate allocations to ensure that every request leads to a safe state. Without knowing these maximum requirements, the system cannot guarantee that it will avoid future deadlock scenarios.

83. (B) Pure Segmentation
Pure segmentation is a memory management technique that divides a process into logical segments of varying sizes. Because these segments must be loaded into contiguous blocks of physical memory, the system often

NODIA

65. (D) 1, 2, 3, 4, and 5
Multithreaded programming introduces several design challenges. Managing process creation with `fork()` and replacing memory with `exec()` requires careful coordination. Signal handling must determine which thread receives notifications. Thread cancellation involves safely stopping execution. Additionally, thread-local storage provides private data for each thread, and scheduler activations help manage communication between the user-level thread library and the operating system kernel.

66. (C) Round Robin (RR)
Round Robin is specifically designed for time-sharing systems where responsiveness is critical for interactive users. Each process is allocated a small, fixed unit of CPU time called a quantum. If a process does not finish within this period, it is preempted and moved to the back of the queue. This ensures that all active processes receive regular turns on the processor.

67. (C) Enforce a total ordering of all resource types and require that each process requests resources in an increasing order.
Circular wait is a condition where a closed chain of processes each waits for a resource held by the next. By establishing a strict numerical order for all resource types and requiring processes to request them in that sequence, the system prevents the possibility of a cycle. If a process needs resource three, it cannot subsequently request resource one, avoiding circular dependencies.

68. (C) Optimal (OPT)
The optimal page-replacement algorithm replaces the page that will not be accessed for the longest duration in the future. While this algorithm provides the lowest possible page-fault rate, it is impossible to implement in practice because it requires perfect knowledge of future memory requests. It serves primarily as a theoretical benchmark for evaluating the performance of other practical replacement algorithms.

69. (B) A track
On a magnetic disk platter, data is organized into concentric circles called tracks. Each platter contains thousands of these tracks. During disk operations, the read/write head stays at a fixed distance from the center to access a specific track as the disk spins. Tracks are further divided into smaller arcs known as sectors, which are the basic units of storage.

70. (C) Worm
A worm is a sophisticated type of system threat that utilizes network connections to spread from one computer to another independently. Unlike viruses, which require a host file to attach to, or Trojan horses, which need user action, worms can self-replicate and exploit system vulnerabilities automatically. This allows them to infect large numbers of connected devices very quickly across the internet.

71. (A) P-1, Q-4, R-3, S-2
File allocation methods define how disk blocks are assigned. Contiguous allocation is fast but causes external fragmentation. Linked allocation connects blocks with pointers, avoiding fragmentation but making random access slow. Indexed allocation uses a dedicated block to store pointers for each file, enabling efficient direct access. The File Allocation Table is a linked variation that stores pointers in a central table.

72. (B) Virtual Memory
Virtual memory is a sophisticated technique that allows a computer to execute programs larger than its physical RAM. The operating system creates an abstraction that gives each process its own large, contiguous address space. Behind the scenes, the system maps this space onto physical memory and secondary storage, moving data back and forth as needed during the execution of programs.

73. (B) A monolithic kernel with modular capabilities.
The Linux kernel is fundamentally monolithic, meaning core services like process management and file systems run in a single address space. However, it incorporates a highly flexible modular design. This allows administrators to dynamically load and unload kernel modules, such as hardware drivers, while the system is running. This approach combines the performance of a monolithic kernel with the flexibility of modularity.

74. (A) Both Statement I and Statement II are correct.
Distributed operating systems aim for high transparency, making a network of separate

Downloaded from www.pdf.tube

across different files. It also performs relocation to ensure that all code and data are correctly addressed within the final program.

93. (A) Both Statement I and Statement II are correct.
Operating systems provide multiple ways for users and programs to interact with the system. Users can use a command-line interface for text-based control, a graphical user interface for visual interaction, or batch systems for automated tasks. For programs, system calls serve as the essential bridge to request services from the kernel, ensuring that hardware resources are managed securely and efficiently.

94. (C) It allows for both concurrency and efficiency by combining multiple models.
The Many-to-Many multithreading model offers a balance between complexity and performance. It multiplexes many user-level threads onto a smaller or equal number of kernel threads. This allows for high concurrency, as multiple threads can run on different processors, and efficiency, as a single blocking system call doesn't necessarily halt the entire process, provided other kernel threads are available to work.

95. (B) Where each processor is self-scheduling and may have its own ready queue.
In symmetric multiprocessing, all processors are treated equally and can perform any task within the operating system. Each processor is responsible for its own scheduling, meaning it can select processes to run from a common ready queue or its own private queue. This architecture improves performance and reliability, as the system can continue functioning even if one processor fails.

96. (C) Preemption
For a deadlock to occur, four conditions must be met simultaneously: mutual exclusion, hold and wait, circular wait, and no preemption. No preemption means that resources cannot be forcibly taken away from a process; they must be released voluntarily. Therefore, preemption itself is a way to resolve or prevent deadlocks, rather than being a necessary condition for them to form.

97. (B) 1, 2, 3, and 4 only
Fundamental memory management schemes include contiguous allocation, swapping, paging, and segmentation. Demand paging is a specific implementation of paging where pages are only loaded as needed. These techniques focus on how processes are assigned physical memory. Thrashing, while related to memory, is a performance problem that occurs when excessive paging activity causes the system to become extremely inefficient.

98. (B) RAID 1
RAID 1, often called mirroring, provides high data reliability by writing the same data to two or more disks simultaneously. It does not use striping, which spreads data across disks, or parity, which uses mathematical checks for recovery. Instead, if one disk fails, the system can immediately continue operating using the mirrored copy, making it ideal for critical system data.

99. (B) To represent disk space as a sequence of bits indicating block status.
A bit vector is a simple and efficient method for managing free space on a disk. Each block is represented by a single bit; for example, a '0' could indicate that the block is free, while a '1' indicates it is currently in use. This allows the operating system to quickly find contiguous blocks of free space by scanning the bit sequence.

NODIA

100. (A) System threat
A system threat involves exploiting vulnerabilities in the operating system's internal mechanisms or services to gain unauthorized access or elevated privileges. Unlike program threats like viruses that infect files, system threats focus on the underlying architecture of the OS itself. Examples include exploiting flaws in memory management or network protocols to bypass security checks and take control of the machine.

101. (C) A network OS is less transparent than a distributed operating system.
The main difference between network and distributed operating systems is the level of transparency. In a distributed OS, the network of computers appears to the user as a single, unified system. In contrast, a network OS is less transparent; users are aware of the separate machines and must explicitly manage and access

Creating an executable from source code involves a series of steps. First, a compiler translates the high-level C++ code into machine-readable object files. Next, a linker combines these object files with required libraries and resolves symbolic references. Finally, a loader takes the resulting executable file and places it into the computer's main memory, preparing it for execution by the central processor.

158. (B) Protecting THE operating system from errant user programs.
Dual-mode operation is a fundamental security feature of modern operating systems. It uses hardware support to distinguish between unprivileged user mode and privileged kernel mode. User applications run in user mode, where they are restricted from accessing hardware directly. If a program needs a privileged service, it must use a system call to switch to kernel mode, allowing the OS to.

159. (A) 1, 4, and 5 only
Deadlock prevention focuses on ensuring that one of the four necessary conditions for deadlock cannot hold. Requiring processes to declare and acquire all resources upfront prevents the hold and wait condition. Enforcing a total ordering of resource types and disallowing circular dependencies structurally ensures that a circular wait can never occur. These proactive strategies make deadlocks impossible within the managed system.

160. (C) Long-term scheduler
The long-term scheduler, also known as the job scheduler, determines which processes are admitted to the system's ready queue. By selecting processes from a job pool and loading them into memory, it controls the degree of multiprogramming. Its goal is to maintain a good mix of I/O-bound and CPU-bound processes to ensure that all system resources are used efficiently.

161. (C) To signal that a process is finished, potentially unblocking others.
The signal operation on a semaphore is used to indicate that a process has released a shared resource or finished a critical section. This operation increments the semaphore's integer value. If other processes were blocked while waiting for that resource, one of them will be moved to the ready queue, allowing it to proceed and ensuring efficient coordination between concurrent tasks.

162. (A) Both A and R are true and R is THE correct explanation of A.
The One-to-One threading model provides great concurrency but can be resource-intensive. Because every user thread requires a corresponding kernel thread, the overhead of creating and managing threads is higher than in other models. Many operating systems limit the total number of threads to prevent the system from becoming overloaded by the sheer number of kernel-level structures and scheduling tasks.

163. (C) An anti-lock braking system in a car.
A real-time operating system is designed for environments where the correctness of a task depends on both the logical result and the time at which it is delivered. In an anti-lock braking system, the OS must respond to sensor data and adjust braking pressure within milliseconds to ensure safety. Standard desktop or batch systems prioritize overall throughput over such strict deadlines.

164. (A) Map logical pages to physical frames.
A page table is a vital data structure used by the memory management unit to implement paging. It stores the mapping between a process's logical pages and the physical frames in RAM. When the CPU generates a logical address, the system uses the page table to find the actual physical location of the data, allowing processes to be stored non-contiguously in memory.

165. (C) FIFO
The First-In First-Out page replacement algorithm is a simple strategy where the operating system tracks the order in which pages are brought into memory. When a new page must be loaded and memory is full, the oldest page is chosen for replacement. While easy to implement, it can be inefficient as it does not consider how frequently a page is used.

166. (D) SSTF (Shortest Seek Time First)
The Shortest Seek Time First algorithm can lead to starvation because it always chooses the request nearest to the current head position. If

Click Here to Buy book on Amazon

111. (A) 1 and 2 only
NTFS is the modern file system used by Windows, offering advanced features like journaling, file-level security through access control lists, and data compression. It uses a Master File Table to store detailed metadata about every file and directory. Unlike the older FAT system, NTFS is designed for reliability and efficiency on large storage volumes, providing robust protection against data corruption.

112. (B) Communication Protocol
A communication protocol is a set of formal rules that define how data is formatted, transmitted, and received across a network. In a distributed system, these protocols are essential for ensuring that different components, often running on diverse hardware and operating systems, can communicate effectively. Examples like TCP/IP provide the necessary structure for reliable data exchange in complex network environments.

113. (C) A compiler translates the entire program before execution, while an interpreter does so line by line.
Compilers and interpreters are two different ways to translate high-level code into a form the CPU can execute. A compiler processes the entire source program at once to produce an independent executable file. An interpreter, on the other hand, translates and executes each instruction one by one during the program's run. Compilers typically offer faster execution, while interpreters provide more flexibility.

114. (C) Kernel
The kernel is the core component of an operating system that remains in physical memory at all times. It is responsible for managing the system's most critical functions, including process scheduling, memory allocation, and hardware communication. As the primary layer between hardware and applications, it ensures that resources are shared fairly and that the system remains stable and secure during operation.

115. (A) 2, 4, 1, 3
CPU scheduling has evolved from simple to complex methods. The oldest is First-Come First-Served, which is easy but inefficient. Shortest Job First was developed to minimize waiting time. Round Robin followed to provide fair time-sharing for interactive users. Multilevel Feedback Queue Scheduling is the most modern and complex, using multiple queues with different priorities to dynamically adapt to process behavior.

116. (B) Ready queue
The ready queue is a data structure maintained by the operating system that holds all processes that are prepared to execute on the processor but are currently waiting for their turn. These processes have all the resources they need, except for the CPU itself. The scheduler selects the next process from this queue based on the specific scheduling algorithm being used.

117. (C) Turn-taking and willingness to enter
Peterson's solution is a classic software-based algorithm that solves the critical-section problem for two processes. It uses two shared variables: an array to indicate a process's interest in entering the critical section and a turn variable to determine whose turn it is. This combination of turn-taking and expressed interest ensures that only one process can enter the critical section at once.

NODIA

118. (B) One-to-One
The One-to-One multithreading model is the most prevalent in modern operating systems. In this model, each user-level thread created by an application is mapped to its own kernel-level thread. This allows for true parallelism on systems with multiple processors and ensures that if one thread blocks, other threads in the same process can continue to execute without interruption.

119. (C) Round Robin
Round Robin is a preemptive scheduling algorithm designed for time-sharing systems. Each process is given a small, fixed unit of time on the CPU, known as a time quantum. If a process does not complete within its quantum, it is forcibly preempted and placed at the back of the ready queue. This ensures that every process gets a regular, fair share.

120. (D) Deadlock Recovery - Granting all resource requests immediately.
Deadlock handling involves several distinct approaches. Prevention and avoidance focus on ensuring a deadlock never occurs by using strict

and provide quick response times for interactive users. This continuous management allows for effective multitasking and ensures that all active processes receive a fair share of processing.

176. (C) No Preemption
No preemption is one of the four necessary conditions for a deadlock. It states that a resource can only be released voluntarily by the process that currently holds it. The operating system cannot forcibly take a resource away to give it to another process. This condition, combined with others like hold and wait, allows for the formation of circular dependencies and deadlocks.

177. (C) Providing a process with more memory frames (its working set).
Thrashing happens when a process spends more time moving pages in and out of memory than executing. To stop this, the operating system must ensure that each process has enough physical memory frames to hold its working set, which is the collection of pages it is currently using. By providing sufficient frames, the page fault rate drops and system performance improves.

178. (C) P-2, Q-1, R-3, S-4
I/O management in an operating system involves several layers. Hardware components like DMA controllers handle direct data transfers. The application I/O interface provides a standard set of system calls for programmers. The kernel subsystem manages efficiency through techniques like spooling and caching. Meanwhile, directory implementation on storage devices often utilizes data structures like hash tables for fast file and metadata.

179. (B) A list for each domain, specifying THE objects and operations it can perform.
A capability list is a row-based implementation of the access matrix security model. It is associated with a domain or user and lists all the objects that the domain is permitted to access, along with the specific rights for each object. In this system, possessing a capability is like having a key that allows access to a specific resource or operation.

180. (C) Hypervisor or Virtual Machine Monitor (VMM)
A hypervisor is a specialized piece of software that enables virtualization by managing multiple virtual machines on a single physical host. It provides each guest operating system with a virtualized version of the hardware, including the CPU, memory, and I/O devices. The hypervisor coordinates resource allocation and ensures that each virtual machine remains isolated and secure from the others on the host.

181. (D) All of THE above
Linux provides a comprehensive set of interprocess communication mechanisms to support diverse application needs. Pipes and FIFOs are used for simple data streams. Message queues allow for the exchange of structured data packets. Semaphores are essential for process synchronization, and shared memory provides high-speed data access. Together, these tools enable complex multi-process applications to coordinate their activities and share data efficiently.

182. (B) Objects with associated attributes and services.
Modern Windows operating systems use an object-oriented approach to manage system resources. Resources like files, processes, threads, and hardware devices are all represented as objects within the kernel. Each object has specific attributes and a set of methods or services that can be used to manipulate it. This design provides a consistent and secure framework for resource management and protection.

183. (A) Both Statement I and Statement II are correct.
Designing a distributed system involves addressing key challenges like robustness and transparency. Robustness ensures that the system can detect and recover from failures in nodes or network links. Transparency aims to hide the complexities of distribution, making the remote resources appear to the user as if they were local. Together, these principles help create reliable and easy-to-use systems across multiple computers.

184. (C) Combine various object modules into a load module.
The primary role of a linker is to take multiple independent object modules, which were produced by a compiler or assembler, and merge them into a single, unified load module or executable. During this process, the linker resolves all symbolic references between the

Downloaded from www.pdf.tube

Programming languages exist at different levels of abstraction. Machine language is the lowest level, consisting of binary code directly understood by hardware. Assembly language adds a small layer of abstraction by using human-readable mnemonics. High-level languages, such as C++ or Python, provide the highest level of abstraction, allowing programmers to write complex logic without worrying about the underlying hardware.

130. (A) Current activity, represented by the value of the program counter.
A process's state is a dynamic property that reflects its current activity during execution. A key part of this state is the program counter, which indicates the address of the next instruction to be performed. Other components of the state include the contents of the CPU registers, the process stack for temporary data, and the data section for global variables.

131. (D) 1, 2, 3, and 4
Process scheduling and operations involve multiple components working together. Long-term schedulers manage system load, while short-term schedulers handle immediate CPU allocation. Process behavior can be I/O-bound or CPU-bound. Context switching is the essential mechanism that saves and restores process states, enabling multitasking and facilitating interprocess communication by allowing the operating system to switch between active tasks during data exchange.

132. (B) Thread cancellation
Thread cancellation is a mechanism used to terminate a thread before it has completed its task. This can happen in two ways: asynchronous cancellation, where the thread is stopped immediately, or deferred cancellation, where the thread periodically checks if it should exit. This is useful in scenarios where a task is no longer needed, such as when a user cancels.

133. (D) OpenMP
OpenMP is a popular example of implicit threading, where the programmer uses compiler directives to specify blocks of code that can be run in parallel. The compiler and the runtime library handle the complexities of thread creation, management, and synchronization. This approach is much simpler than explicit threading, where the developer must manually manage every aspect of the thread's lifecycle.

134. (A) Both A and R are true and R is the correct explanation of A.
Priority scheduling can lead to starvation, where low-priority processes never get to run because high-priority tasks keep arriving. Aging is an effective solution to this problem. It involves gradually increasing the priority of a process the longer it waits in the ready queue. Eventually, even the lowest-priority process will gain a high enough priority to be selected for execution.

135. (C) Terminate the process that has consumed the least amount of CPU time.
When recovering from a deadlock through process termination, the goal is to resolve the situation with the least amount of wasted work. Choosing a victim process that has only recently started and consumed very little CPU time is a common strategy. This minimizes the impact on overall system progress compared to terminating a long-running process that is nearly finished.

136. (B) Logical address
A logical address is a virtual address generated by the CPU during the execution of a program. It is part of the process's private address space and does not correspond directly to a physical location in RAM. The memory management unit is responsible for translating these logical addresses into physical addresses, allowing the system to manage memory flexibly and securely.

137. (B) First-In, First-Out (FIFO)
Belady's Anomaly is a counter-intuitive phenomenon observed in some page-replacement algorithms, most notably First-In First-Out. It occurs when increasing the number of physical memory frames allocated to a process actually results in an increase in the number of page faults. This happens because FIFO does not account for the frequency or timing of page access, leading to inefficient memory usage.

138. (B) Track
As a magnetic disk platter rotates, the area that passes directly under the stationary read/write head forms a circular path known as a track. A typical disk surface has thousands of these tracks.

NODIA

rules or predictive algorithms. Detection identifies if a deadlock has happened. Recovery, however, is the process of breaking an already existing deadlock. This usually involves terminating one or more processes or preempting resources, rather than granting even more requests which would only worsen.

121. (D) P-3, Q-4, R-2, S-1
Memory management techniques fulfill different roles. Demand paging loads data only when needed, while page replacement decides which frame to free when memory is full. Thrashing is the performance failure caused by excessive paging. Segmentation organizes memory into logical, variable-sized units. Together, these concepts allow the operating system to manage physical RAM efficiently while supporting large, complex applications and multitasking.

122. (A) A file on disk to be mapped to the virtual address space for routine memory access.
Memory-mapped files are a powerful feature that allows the operating system to map a file's content directly into a process's virtual address space. This enables the program to access the file's data using standard memory operations, such as pointers, instead of traditional read and write system calls. This can significantly improve performance for applications that frequently access large data files.

123. (B) Rotational latency
Rotational latency is the time it takes for the specific sector of data on a spinning disk platter to reach the read/write head. This delay depends on the rotation speed of the disk, measured in revolutions per minute. Along with seek time, rotational latency is a major component of the total disk access time, impacting the performance of storage operations.

124. (A) THE ability for multiple users to access the same file concurrently with appropriate protection.
File sharing is a critical feature of multi-user operating systems, allowing several people to access and collaborate on the same data. To ensure security and data integrity, the system implements protection mechanisms, such as permissions and locks. These controls define which users can read, write, or execute a file, preventing unauthorized access while enabling efficient information exchange in shared environments.

125. (C) Denial of Service (DoS)
Program threats are malicious software that infect or exploit individual applications, such as Trojan horses, trap doors, and logic bombs. A Denial of Service attack, however, is classified as a system or network threat. Its goal is to make a computer or network resource unavailable to its intended users by flooding it with an overwhelming amount of traffic or requests.

126. (D) Statement I is incorrect but Statement II is correct.
While the Linux kernel is monolithic, meaning its main components run in a single address space, it is not a single, unchangeable executable. Its design is enhanced by loadable kernel modules, which allow new functionality, such as hardware drivers, to be added or removed while the system is running. This modularity provides great flexibility without sacrificing the performance of a monolithic.

127. (B) A layered architecture with a Hardware Abstraction Layer (HAL) for portability.
The Windows NT family, which includes all modern versions of Windows, is designed with a layered architecture for maximum portability and stability. A key component is the Hardware Abstraction Layer, which sits between the physical hardware and the rest of the OS. This design allows the same core operating system to run on different hardware architectures with minimal changes to the software.

128. (B) Continue operating correctly despite failures of some of its components.
Robustness is a vital characteristic of a distributed system, referring to its ability to handle failures gracefully. Because these systems consist of many interconnected machines, the chance of a single component failing is relatively high. A robust system uses redundancy and error-recovery techniques to ensure that the overall service remains available and correct even when specific nodes or links fail.

129. (A) 2, 1, 3

Click Here to Buy book on Amazon

improve performance by reducing cache misses and expensive memory reloads.

148. (C) Require processes to request resources in a predefined order.
Deadlock prevention involves structurally ensuring that at least one of the four necessary conditions for deadlock cannot occur. By requiring processes to request all resources in a strict, predefined numerical order, the circular wait condition is eliminated. This ensures that no cycle can ever form in the resource-allocation graph, thereby guaranteeing that the system remains free from deadlocks during its operation.

149. (A) P-4, Q-1, R-2, S-3
Security threats come in many forms. A Trojan horse is a malicious program disguised as legitimate. A worm is a self-replicating program that spreads across networks. Denial of Service attacks aim to crash or flood a system to prevent normal use. A trap door is a hidden entry point in software that allows unauthorized access by bypassing normal security protocols and checks.

150. (B) It can suffer from significant external fragmentation.
Pure segmentation allows a process's memory to be divided into logical units of different sizes. However, as segments are loaded and removed from memory, the free space becomes broken into many small, non-contiguous holes. This external fragmentation can make it impossible to load a new segment, even if the total amount of free memory is more than enough for the request.

151. (A) Both Statement I and Statement II are correct.
Contiguous file allocation offers excellent read performance because the data blocks are stored sequentially on the disk, minimizing head movement. However, it suffers from the significant drawback of external fragmentation. As files are deleted and created, it becomes difficult to find a single, large enough block of contiguous space for new files, leading to inefficient use of the storage device.

152. (A) RAID 0
RAID 0, also known as striping, is designed for maximum performance rather than data safety. It spreads data across multiple disks, allowing for faster parallel read and write operations. However, it provides no redundancy; if any single disk in the array fails, all the data in the entire array is lost. This makes it unsuitable for storing critical or irreplaceable information.

153. (C) Holds data for a device in a buffer on disk until THE device is ready.
Spooling is a technique used to manage I/O for slow devices, such as printers. Instead of the CPU waiting for the device to finish, the data is written to a buffer on the disk. The operating system then sends the data from the disk to the device at its own speed, allowing the CPU to continue with other tasks.

154. (B) Taking away previously granted permissions from a user or process.
Revocation is a vital part of access control, involving the removal of permissions from a user or a domain. This can be challenging to implement efficiently, especially in large systems. Common strategies include searching all access control lists or using indirect pointers that can be invalidated. Revocation ensures that security remains up-to-date as users change roles or leave an organization.

155. (D) A fixed, 8.3 filename limitation
NTFS is a sophisticated file system used by Windows that supports many modern features, including journaling for reliability, very large file and volume sizes, and detailed security through access control lists. It does not have the restrictive 8.3 filename limitation of the older FAT file systems. Instead, NTFS allows for long, descriptive filenames, making it much easier for users to organize.

156. (B) To provide a standard way for machines on a network to exchange data.
A communication protocol like TCP/IP is a set of rules that standardizes how data is sent across a network. In a distributed system, these protocols are essential for ensuring that diverse computers can connect and share information reliably. They handle tasks like addressing, routing, and error checking, providing a stable foundation for complex network-based applications and services to operate across machines.

157. (A) Compilation, Linking, Loading

Downloaded from www.pdf.tube

Data is recorded onto these tracks in small units called sectors. To access different data, the disk arm must move the head to the correct track on the platter.

139. (C) Both a linear list and a hash table are possible implementations.
Directory implementation focuses on how to store and look up filenames on a disk. A linear list is the simplest method, but searching through it can be slow for large directories. A hash table is much more efficient, allowing for nearly instantaneous lookup by using a hash function on the filename. Many modern file systems use a combination of these techniques.

140. (D) Maximizing the speed of I/O operations.
Protection in a computer system is focused on security and resource management. Its primary goals are to prevent processes from interfering with each other, ensure that resources are accessed only by authorized users, and enforce system-wide policies for resource usage. While efficient I/O is important for performance, it is not a goal of the protection mechanisms themselves, which focus on safety.

141. (B) Kernel Modules
Loadable kernel modules are a key feature of the Linux operating system, allowing it to remain flexible and efficient. These modules can be added to or removed from the running kernel on demand, without needing to reboot the system. This allows for the dynamic addition of hardware drivers, file system support, and other system functions, keeping the core kernel lightweight and adaptable.

142. (A) Location transparency
Location transparency is a fundamental design goal of a Distributed File System. It ensures that a user can access a file using its name without needing to know the physical location of the server where the file is stored. This makes the distributed system easier to use and allows administrators to move files between servers for load balancing without affecting users.

143. (C) Loading is the process of resolving references between different object files.
System software tools have distinct roles in the software lifecycle. A macro is expanded into code, while a debugger helps find errors. An interpreter executes code statement by statement. The process of resolving references between different object files and combining them into a single file is actually called linking. Loading is the subsequent step where the final executable is placed into memory.

144. (C) ROM or EEPROM (firmware)
The bootstrap program is the very first piece of software that runs when a computer is turned on. Because physical RAM is empty at startup, this program must be stored in non-volatile memory, such as Read-Only Memory or Electrically Erasable Programmable ROM. This firmware initializes the hardware and then locates and loads the operating system kernel from the disk.

145. (B) 1, 2, and 3 only
Solving the critical-section problem requires meeting three criteria: mutual exclusion, progress, and bounded waiting. Semaphores are a key software-based synchronization tool, and spinlocks are characterized by busy waiting. While Peterson's solution is a classic software method, it is mathematically proven to work only for two processes, although it provides an important foundation for more complex synchronization algorithms in modern systems.

146. (C) A blocking system call by one thread blocks the entire process.
The Many-to-One multithreading model is inefficient for modern multitasking. Because multiple user-level threads are mapped to only one kernel-level thread, the operating system sees the entire process as a single unit. If any one thread in the process performs a blocking operation, such as waiting for I/O, the entire process is put to sleep, preventing all other threads.

147. (A) THE tendency for a process to be assigned to the same processor it was previously running on.
Processor affinity is a scheduling concept in multiprocessor systems where the operating system tries to keep a process running on the same CPU core. This is beneficial because the processor's local cache likely still contains data from that process. By avoiding moving the process to a different core, the system can

NODIA

a steady stream of new requests arrives near the head, the disk arm may stay in one area for a long time. Requests located far from the current head position might be indefinitely postponed, waiting for a break in the local traffic.

167. (B) An object, specifying which users/processes have what access rights to it.
An access control list is a security mechanism associated with an object, such as a file or directory. It contains a list of entries that specify which users or processes are granted access and what specific operations they can perform, such as reading, writing, or executing. This allows for fine-grained control over resource security in a multi-user operating system environment.

168. (C) Create a child process and have it run a new program.
In UNIX and Linux, creating a new program involves two distinct steps. First, the `fork()` system call creates a duplicate of the current process, resulting in a parent and a child. Then, the child process typically calls `exec()`, which replaces its entire memory space and code with a new executable file, allowing it to run a completely different program.

169. (B) A set of rules that govern THE format and transmission of data.
A protocol is a formal set of rules and conventions that determine how data is exchanged between electronic devices. In network communication, protocols define the structure of data packets, the methods for error detection and correction, and the procedures for initiating and terminating a connection. This standardization is essential for ensuring that different computers and devices can communicate reliably and efficiently.

170. (C) Java
Java is a prime example of a high-level programming language designed for portability and ease of use. It uses human-readable syntax and abstracts away complex hardware details, such as memory management. Unlike low-level assembly or machine code, which are specific to a particular CPU architecture, Java programs are typically compiled into bytecode and run on a virtual machine, ensuring broad compatibility.

171. (C) THE contents of THE user stack
A Process Control Block is a data structure used by the operating system to store all the information needed to manage a specific process. This includes the process state, program counter, CPU registers, and scheduling information. While it tracks the process's memory management details, the actual content of the user stack, which contains temporary function data, is stored in physical memory.

172. (A) Shared memory IPC is generally faster than message passing IPC.
Shared memory is an efficient interprocess communication mechanism because it allows multiple processes to access the same region of physical memory directly. This avoids the overhead of copying data between process address spaces, which is necessary in message-passing systems. However, shared memory requires careful synchronization using tools like semaphores to prevent multiple processes from causing data corruption during concurrent access.

173. (A) 3, 1, 4, 5, 2
When a page fault occurs, the process first traps to the operating system. The OS then searches for a free frame in physical memory. Once a frame is found, the required page is loaded from the disk into that frame. The page table is then updated to reflect the new mapping. Finally, the process is restarted at the exact instruction that caused.

174. (A) Determining which thread in a process should receive a signal.
Signal handling in multithreaded environments is a complex design issue. When a signal is generated for a process, the operating system must decide which specific thread should receive it. Options include delivering it only to the thread it applies to, to all threads in the process, or to a specific designated signal-handling thread. This choice significantly impacts the behavior and reliability.

175. (A) To keep THE CPU as busy as possible and provide fast response times to users.
The primary goal of CPU scheduling is to maximize system efficiency by ensuring the processor is rarely idle. In a multiprogramming environment, the scheduler selects the next process for execution to maintain high throughput

Click Here to Buy book on Amazon

resources across different nodes using specific network commands.

102. (C) Metadata such as owner, permissions, and pointers to data blocks.
An inode is a data structure in UNIX and Linux file systems that stores essential metadata about a file. It contains information such as the file's owner, access permissions, timestamps, and size. Most importantly, it holds pointers to the actual disk blocks where the file's data is stored. Crucially, the inode does not store the filename itself.

103. (C) Semaphore - A variable used for signaling among processes.
A semaphore is a fundamental synchronization tool used to coordinate the activities of multiple processes. It is an integer variable that is accessed through two atomic operations: wait and signal. These operations allow processes to signal each other about the state of shared resources or events, helping to prevent race conditions and manage access to critical sections of code.

104. (C) Disk Bandwidth
CPU scheduling criteria are metrics used to evaluate the efficiency of scheduling algorithms. Throughput measures the number of processes completed per unit of time, waiting time is the total time spent in the ready queue, and response time is the delay before the first response. Disk bandwidth, however, is a measure of data transfer speed for storage devices, not CPU.

105. (A) Both A and R are true and R is the correct explanation of A.
In systems where each resource type has only a single instance, a cycle in the wait-for graph is both a necessary and sufficient condition for a deadlock. The wait-for graph is a simplified version of the resource-allocation graph that only shows the dependencies between processes. If a cycle exists, it means a set of processes is permanently blocked.

106. (B) A page is logical memory, while a frame is a block of physical memory.
In paging systems, memory is viewed in two different ways. The logical address space of a process is divided into fixed-size units called pages. Physical memory is divided into units of the same size called frames. The operating system maintains a page table to map each logical page to a corresponding physical frame, allowing processes to be stored non-contiguously in RAM.

107. (C) SCAN (Elevator)
The SCAN algorithm, also known as the elevator algorithm, is designed to be fair and efficient. The disk head moves from one end of the disk to the other, servicing all requests along the way. When it reaches the end, it reverses direction and repeats the process. This bidirectionality ensures that all requests are eventually serviced, preventing the starvation of any.

108. (A) A set of system calls.
The Application I/O Interface provides a standardized way for user programs to interact with the diverse range of hardware devices connected to a computer. This is implemented through a specific set of system calls, such as open, read, write, and close. These calls abstract the hardware complexities, allowing developers to write software that can communicate with many different devices.

109. (B) To ensure that a user is who they claim to be before granting access to the system.
User authentication is the first line of defense in computer security. Its primary goal is to verify the identity of an individual trying to access the system. By using passwords, biometrics, or security tokens, the operating system ensures that only authorized users can log in. Once authenticated, the system can then apply authorization rules to control what resources they access.

110. (B) The creation of a virtual (rather than actual) version of something, such as a hardware platform, operating system, or storage device.
Virtualization is a powerful technology that creates an abstraction layer between physical hardware and the software running on it. It allows a single physical machine to be partitioned into multiple virtual machines, each running its own independent operating system and applications. This technology is fundamental to cloud computing, enabling efficient resource sharing, improved scalability, and simplified management of complex computer systems.

NODIA

modules, ensuring that function calls and variable accesses correctly point to their intended definitions within the final program.

185. (B) A program that acts as an intermediary between THE user and hardware.
An operating system is a fundamental piece of software that manages a computer's hardware and provides a platform for applications to run. It acts as an intermediary, abstracting the hardware's complexity and providing users with a more convenient interface. By controlling resource allocation and providing common services, the operating system ensures that the computer is used efficiently, securely, and reliably.

186. (C) 1, 2, 3, 4, and 5
A typical process life cycle involves several distinct states. A process starts in the 'New' state during creation, then moves to 'Ready' while waiting for CPU time. Once it begins executing, it is 'Running'. If it needs to wait for an event, like I/O, it enters the 'Waiting' state. Finally, after completing its tasks, the process reaches the 'Terminated' state.

187. (A) Mutual Exclusion, Progress, Bounded Waiting
A valid solution to the critical-section problem must satisfy three essential requirements. Mutual exclusion ensures that only one process can be in its critical section at a time. Progress means that if no process is in its critical section, the selection of the next process cannot be delayed indefinitely. Bounded waiting guarantees that every process will eventually be allowed to enter its.

188. (B) A standard IEEE specification for thread behavior.
Pthreads, short for POSIX Threads, is a standardized model for creating and managing threads defined by the IEEE. It provides a common set of C language programming interfaces for thread creation, synchronization, and management. This standard ensures that multithreaded code is portable across different UNIX-like operating systems, allowing developers to write concurrent applications that run on various platforms.

189. (B) SJF (Shortest Job First)
The Shortest Job First scheduling algorithm is mathematically proven to provide the minimum average waiting time for a given set of processes. By prioritizing the tasks with the smallest CPU burst times, it allows more processes to finish quickly, reducing the overall time others spend waiting in the ready queue. This makes it an ideal theoretical benchmark for CPU efficiency.

190. (B) Both A and R are true but R is not THE correct explanation of A.
Deadlock avoidance is a proactive strategy that uses predictive information to ensure the system never enters an unsafe state. Algorithms like the Banker's algorithm require each process to declare its maximum resource needs in advance. This allows the system to simulate each allocation and only grant requests if it can guarantee that a safe sequence of process completion still exists.

191. (A) Translating logical addresses to physical addresses.
The memory management unit is a hardware component that handles the mapping of virtual memory to physical RAM. Its primary task is to translate the logical addresses generated by the CPU into the physical addresses used by the memory controller. This translation allows the operating system to implement features like paging, memory protection, and address space isolation between different active processes.

192. (D) Seek Time - THE time for THE desired sector to rotate under THE head.
Disk performance is determined by several physical factors. Tracks are concentric paths on a platter, and sectors are the basic units of data on those tracks. A cylinder represents all tracks at the same arm position. Seek time is actually the time required to move the disk arm to the correct track, not the time spent waiting for a sector to rotate.

193. (B) Linked allocation
Linked allocation is a file storage method where the data blocks of a file can be scattered anywhere on the disk. Each block contains a pointer to the next block in the sequence, forming a linked list. This method eliminates external fragmentation and allows files to grow easily, but it makes random access very slow because the pointers must be followed sequentially.

Downloaded from www.pdf.tube

where processes are running on different computers. The operating system handles the complexities of transferring the messages across the network, providing a reliable way for distant tasks.

213. (A) 3, 1, 4, 2
Storage devices form a hierarchy based on access speed and cost. CPU registers are the fastest, providing immediate access for the processor. Main memory is also very fast but much larger. Solid-state disks are a slower but high-capacity form of non-volatile storage. Magnetic disks are the slowest due to their physical moving parts, such as spinning platters and moving actuator arms.

214. (D) Kernel-level threads are invisible to THE operating system.
Threads are managed differently depending on the model. User-level threads are faster to manage because they don't require kernel involvement. The One-to-One model provides good concurrency by mapping user threads to kernel threads. Crucially, kernel-level threads are fully managed by and visible to the operating system, which is responsible for scheduling them and handling their system calls and I/O operations.

215. (B) Starvation
Starvation is a classic problem in priority scheduling where a low-priority process is indefinitely delayed because higher-priority tasks are always available to run. The process remains in the ready queue but never receives CPU time. This can happen in preemptive systems if there is a continuous stream of important jobs, necessitating solutions like aging to eventually elevate the process's priority.

216. (D) Augmentation
Deadlock handling involves several recognized strategies to ensure system stability. Prevention ensures deadlocks cannot occur by negating necessary conditions. Avoidance uses predictive algorithms like the Banker's algorithm. Detection and recovery identify and resolve existing deadlocks. Augmentation is not a standard term or method in the context of deadlock management, which focuses on these three primary approaches for maintaining efficient and stalled-free operation.

217. (A) THE page fault rate and THE time to service a page fault.
In a paging system, the effective memory access time is the average time taken to retrieve data. This time is heavily influenced by how often page faults occur. When a fault happens, the system must wait for the disk to load the page, which is very slow. Therefore, even a small increase in the page fault rate can significantly impact.

218. (B) Directory or Volume Table of Contents
A directory, or Volume Table of Contents, is a vital file system structure that keeps track of all the files stored on a specific volume. It contains the names of the files and pointers to their metadata, such as file control blocks or inodes. This organization allows the operating system to locate and manage files efficiently in a large storage space.

NODIA

219. (A) Both A and R are true and R is THE correct explanation of A.
Cryptography is a cornerstone of computer security, providing essential tools for protecting information. By transforming plaintext into an unreadable ciphertext, it ensures the confidentiality of data during storage or transmission. It also provides mechanisms for verifying the integrity of data and the identity of users. Together, these features make cryptography a fundamental tool for implementing robust defenses against unauthorized access.

220. (B) Type 2 (hosted)
A Type 2 virtual machine, also known as a hosted hypervisor, runs as an application layer on top of a standard host operating system. It relies on the host OS to manage the physical hardware and provides a virtual environment for the guest operating system. This type is common on desktop computers, allowing users to run different systems for testing or specialized.

221. (C) P-3, Q-2, R-1, S-4
Linux architecture is built on several key components. Process management uses `fork()` and `exec()` for task creation. The file system often employs ext4 for reliable data storage. Scheduling is handled by the Completely Fair Scheduler

the process of making an entire file system on a storage device available at a specific directory mount point.

204. (B) Allow devices to transfer data to memory without CPU.
Direct Memory Access is a technique used to improve system performance by offloading data transfer tasks from the CPU. It allows high-speed I/O devices, like disk controllers or network cards, to transfer large blocks of data directly to or from main memory. By bypassing the CPU during the transfer, DMA frees the processor to perform other important computational tasks simultaneously.

205. (B) Authentication
Authentication is the process of verifying a user's identity before allowing them to access a computer system. It typically involves checking something the user knows, has, or is. This step is essential for establishing trust and ensuring that only authorized individuals can gain access. Once a user is authenticated, the system then uses authorization rules to determine their specific rights.

206. (A) Both Statement I and Statement II are correct.
Virtualization comes in different forms to meet various performance and compatibility needs. Para-virtualization requires the guest operating system to be modified so it can communicate directly with the hypervisor for better efficiency. Full virtualization, on the other hand, allows an unmodified guest OS to run in complete isolation, as the hypervisor provides a perfect virtual representation of the underlying hardware.

207. (B) To allocate CPU time to runnable processes according to a scheduling policy.
The scheduling subsystem in Linux is responsible for deciding which of the ready processes should be executed by the CPU at any given moment. It uses a specific scheduling policy, such as the Completely Fair Scheduler, to ensure that all tasks receive a fair share of processor time. This continuous management is vital for maintaining high performance and system responsiveness.

208. (B) Executive - Manages core OS services like memory and I/O.
The Windows architecture is composed of several key layers. The Executive is a group of kernel-mode components that provide vital services such as process management, memory allocation, security, and input/output control. It sits on top of the kernel and the Hardware Abstraction Layer. This structure ensures that core operating system functions are handled efficiently and securely within a privileged execution environment.

209. (B) Handle more work by adding resources to THE system.
Scalability is a critical design issue for distributed systems, referring to the ability to handle an increasing workload by adding more resources. This can involve adding more computers to the network or upgrading existing ones. A scalable system maintains its performance and efficiency as it grows, making it suitable for large-scale applications that must support many concurrent users and processes.

210. (C) A set of binary-coded instructions that can be directly executed by a computer's CPU.
Machine language is the most basic level of programming, consisting of instructions encoded as sequences of binary digits. It is the only language that a computer's hardware can understand and execute directly. Because it is highly specific to a particular CPU architecture and very difficult for humans to read, programmers typically use higher-level languages that are then translated into.

211. (B) Microkernel
A microkernel architecture moves many traditional operating system services, such as file systems and drivers, from the kernel into user space. While this improves modularity and security, it often results in performance overhead. This is because every service request requires a context switch from user mode to kernel mode and back, which can be a time-consuming operation for the processor.

212. (C) It is useful for communication across different computers.
Message passing is an interprocess communication mechanism where processes exchange data by sending and receiving messages. Unlike shared memory, it does not require a common address space, making it ideal for distributed systems

Click Here to Buy book on Amazon

Within a single process, multiple threads share several key operating system resources to improve efficiency. While each thread maintains its own unique identifier, program counter, and stack, they all operate within the same code and data sections. This shared environment allows threads to communicate easily and utilize the same set of global variables and open files very efficiently.

267. (B) Implicit threading
Implicit threading is a modern programming approach where the responsibility for creating and managing threads is transferred from the application developer to compilers and runtime libraries. By automating these complex tasks, developers can focus on the core logic of their applications while the system handles the underlying details of parallel execution, load balancing, and thread lifecycle management automatically for them.

268. (A) To map user-level threads to kernel-level threads.
Multithreading models establish the relationship between threads created at the user level by programmers and the threads managed by the operating system kernel. These models, such as many-to-one or one-to-one, define how user threads are mapped onto available kernel threads. This mapping is essential for the operating system to schedule thread execution on the physical processor cores available.

269. (C) The process of thread cancellation and signal handling.
Developing multithreaded applications introduces specific challenges that do not exist in single-threaded environments. Key issues include managing thread cancellation, where one thread terminates another, and signal handling, which determines how asynchronous events are delivered to specific threads. Additionally, maintaining thread-specific data requires careful architectural planning to ensure that parallel execution remains both correct and highly efficient at runtime.

270. (D) The fundamental unit of CPU utilization.
A thread serves as the smallest and most fundamental unit of CPU utilization within a computer system. It represents an independent path of execution that operates within the context of a larger process. By allowing multiple threads to run concurrently, systems can achieve higher levels of parallelism, especially on modern multicore processors, thereby improving overall application performance for users.

271. (C) Turnaround time.
Turnaround time is a vital scheduling criterion that measures the total elapsed time from the moment a process is submitted to the system until its final completion. This metric includes the time spent waiting to enter memory, waiting in the ready queue, executing on the CPU, and performing any necessary input or output operations during its entire lifecycle.

272. (B) Each processor is self-scheduling and takes processes from a common ready queue.
Symmetric multiprocessing is a multi-processor scheduling approach where each individual processor is responsible for its own scheduling decisions. In this model, every CPU examines the common ready queue to select the next process or thread for execution. This decentralized design eliminates the bottleneck of a single master processor, providing better scalability and performance in modern computer systems today.

273. (D) Hard real-time systems must meet task deadlines, whereas soft real-time systems only aim to.
Real-time systems are categorized based on their ability to guarantee task completion within specific time constraints. Hard real-time systems are designed with the strict requirement that every critical task must be serviced by its deadline. In contrast, soft real-time systems prioritize important processes and aim to meet deadlines, but they do not provide an absolute guarantee of timely execution.

274. (A) Average waiting time
The Shortest-Job-First scheduling algorithm is mathematically proven to be optimal for minimizing the average waiting time of a set of processes. By selecting the task with the smallest next CPU burst, it reduces the amount of time that shorter processes must wait behind longer ones. However, implementing this algorithm effectively requires accurate future knowledge of burst lengths.

NODIA

to ensure balanced CPU time. Interprocess communication utilizes pipes and sockets for data exchange, providing a robust framework for complex applications to coordinate effectively.

222. (B) A client-server model for internal operations, with a microkernel-like base.
Modern Windows versions are designed with a flexible architecture that uses a client-server model for many internal operations. It features a relatively small kernel that provides core services, while other OS functions run as separate processes. This design, influenced by microkernel principles, improves system modularity and stability, as a failure in one service process is less likely to crash the entire system.

223. (B) Each computer on THE network runs its own OS and users are aware of different machines.
A network based operating system connects multiple autonomous computers. Unlike a distributed system, which aims to hide the network's existence, a network OS provides less transparency. Users are explicitly aware of the different machines on the network and must manually manage file transfers and remote access. Each computer maintains its own local operating system and handles its own users and resources.

224. (C) Linker or Loader
Relocation is the process of adjusting the memory addresses in a program so that it can be correctly executed from a specific location in RAM. This task is typically performed by the linker when combining object files or by the loader when placing the program into memory. It ensures that all symbolic references and jump instructions point to the correct physical addresses during.

225. (B) THE operating system
A system call is the essential bridge between a running application and the operating system. It provides a programmatic way for software to request high-level services from the kernel, such as file management, process control, and network communication. This interface ensures that hardware resources are accessed in a secure and controlled manner, maintaining the overall stability and integrity of the computer system.

226. (D) 1, 2, 3, 4, and 5
Process management is a core function of the operating system, encompassing a wide range of responsibilities. These include creating and deleting processes, managing their suspension and resumption, and providing essential mechanisms for synchronization and communication. Furthermore, the operating system must implement strategies for deadlock handling to ensure that the system remains responsive and efficient even when multiple processes compete for resources.

NODIA

227. (B) Enters a loop, repeatedly checking if THE lock is available.
A spinlock is a low-level synchronization mechanism where a thread trying to acquire a lock will continuously spin in a loop, repeatedly checking the lock's status. While this avoids the overhead of putting the thread to sleep and performing a context switch, it consumes CPU cycles. Spinlocks are most effective in multiprocessor systems where locks are held for very short durations.

228. (C) Responsiveness, by allowing a program to continue running even if part of it is blocked.
Multithreading significantly improves the responsiveness of an application by allowing different parts of the program to run concurrently. For example, a web browser can use one thread to handle user input while another thread downloads data in the background. This ensures that the application remains interactive and doesn't appear frozen to the user, even when performing long or potentially blocking operations.

229. (C) First-Come, First-Served (FCFS)
Non-preemptive scheduling algorithms are those where once a process starts executing on the CPU, it keeps the processor until it either terminates or blocks for I/O. First-Come First-Served is a classic example of this approach. It is simple to implement but can lead to long waiting times for short processes if a very long task is at the front.

230. (C) No deadlock exists in THE system.
A resource-allocation graph is a powerful tool for analyzing the state of a system. If the graph contains no cycles, it is mathematically guaranteed that no deadlock currently exists. Every process in the system can eventually

249. (C) It limits THE maximum size of THE file.
In indexed allocation, each file has an index block that contains pointers to its data blocks. If this index block is too small, it can only hold a limited number of pointers, which directly restricts the maximum size a file can reach. To overcome this, systems use techniques like multi-level indexing or linked index blocks to support much larger files.

250. (A) Both Statement I and Statement II are correct.
Virtualization is managed by hypervisors, which come in two primary types. A Type 1 hypervisor, also known as bare-metal, is installed directly on the physical hardware and provides excellent performance. A Type 2 hypervisor runs as an application on a host operating system, making it easier to install and use on standard desktop computers for testing and various development.

251. (B) A compiler translates the entire source program at once before execution, while an interpreter translates and executes code sequentially.
A compiler processes the complete source code to create an executable file before any part of the program runs. In contrast, an interpreter processes the source code sequentially, translating and executing one instruction at a time. This fundamental difference in the translation timing and method defines how these tools prepare high-level code for the CPU for efficient software execution.

252. (A) To combine various object modules and resolve external references into a single executable.
The linking process is a critical stage in software development that integrates multiple separate object modules generated during compilation. Its primary responsibility involves resolving external references between these modules to form a single, unified executable file. This enables the loader to subsequently place the complete program into memory for final execution by the computer system during the runtime phase.

253. (C) Relocation
Relocation is the specific task performed by a loader to adjust internal program addresses so they correspond to the actual physical memory locations assigned during execution. Since programs are often compiled with relative addresses, this step ensures that all memory references within the code correctly point to the intended data and instructions once loaded into RAM for active processing.

254. (D) Replacing a shorthand notation with a predefined block of code before compilation.
Macros serve as a productivity tool used during the pre-processing phase of software development. They function by identifying specific shorthand notations or identifiers within the source code and automatically replacing them with a predefined, larger sequence of statements. This mechanism allows developers to simplify complex code patterns and reduce repetitive typing before the actual compilation process begins for the program.

255. (C) Debugger
A debugger is an essential utility designed to assist programmers in the post-execution analysis and error correction of software. It provides the necessary functionality to control the flow of program execution, inspect the current state of variables and memory, and identify specific logic or runtime faults. These tools are indispensable for maintaining software reliability and correctness throughout development cycles.

256. (B) To act as an intermediary that manages hardware and software resources for user applications.
The operating system functions as a fundamental layer of software that bridges the gap between the computer hardware and user-level applications. It acts as an intermediary, managing all hardware and software resources to provide a stable platform for program execution. This centralized control ensures that resources like memory and processing power are efficiently allocated and utilized by the system.

257. (D) By invoking system calls through a defined interface.
User applications are restricted from accessing computer hardware directly to maintain system integrity. Instead, they must interact with the operating system kernel through a standardized interface known as system calls. These calls serve as a formal request mechanism, allowing applications to securely utilize essential services

Downloaded from www.pdf.tube

240. (B) A critical real-time task gets priority over other tasks, but THE deadline is not mandatory.
In a soft real-time system, critical tasks are given priority over regular processes to ensure they are executed as quickly as possible. However, unlike hard real-time systems, there is no absolute guarantee that a task will meet its deadline. Missing a deadline in a soft real-time system may result in degraded performance, but it does not cause a total system failure.

241. (A) Both A and R are true and R is THE correct explanation of A.
In systems with multiple instances of a resource type, a cycle in the resource-allocation graph does not always mean a deadlock. This is because a process within the cycle might be waiting for a resource that could be released by a process outside the cycle. If that external process finishes and frees its resources, the cycle can be broken, allowing others to proceed.

242. (A) THE anomaly where THE page fault rate increases as THE number of allocated frames increases.
Belady's Anomaly is a surprising behavior found in some page-replacement algorithms, such as First-In First-Out. It occurs when a process is given more physical memory frames, but the number of page faults actually goes up instead of down. This counter-intuitive result happens because these algorithms don't account for the frequency of use, leading to inefficient management of the additional memory resources.

243. (B) Selects THE request with THE minimum seek time from THE current head position.
The Shortest Seek Time First disk scheduling algorithm is designed to improve performance by minimizing the movement of the disk head. It always selects the pending request that is geographically closest to the current head position. While this significantly reduces the average seek time compared to FCFS, it can unfortunately lead to the starvation of requests that are located far away.

244. (C) THE attributes and disk block locations of THE file's data.
In UNIX and Linux file systems, an inode is a crucial data structure that stores all the metadata for a file, except for its name. It contains attributes like permissions, owner, and size, and provides the addresses of the physical disk blocks where the file's data is stored. This information allows the operating system to manage and access files efficiently and securely.

245. (B) THE file name does not reveal THE file's storage location.
Location transparency is a key design principle for distributed file systems. It ensures that the name of a file is independent of its physical storage location on the network. This allows users to access files without needing to know which server they reside on, making the system easier to use and allowing administrators to move data between servers without affecting user access.

246. (C) THE bootstrap loader
The bootstrap loader is a small, specialized program stored in non-volatile memory that initiates the boot process. Its primary responsibility is to locate the operating system kernel on the disk and load it into the main memory. Once the kernel is loaded and running, it takes over the management of the computer system and starts the various services and background processes.

NODIA

247. (B) If no process is in its critical section, selection of THE next cannot be postponed indefinitely.
The progress requirement is a vital part of the solution to the critical-section problem. It ensures that if no process is currently executing in its critical section and some processes want to enter, the system must choose one of them to proceed. This choice cannot be delayed forever, ensuring that the system as a whole continues to move forward and perform work.

248. (C) Threads share memory and resources, making communication and sharing easier.
A major advantage of threads over processes is that all threads within a single process share the same address space and system resources. This sharing makes communication between threads much faster and simpler than interprocess communication, which requires complex mechanisms like message passing or shared memory. It allows concurrent tasks to easily access and modify shared data structures without significant operating system.

194. (B) A hidden entry point in a program that bypasses normal security checks.
A trap door, also known as a back door, is a secret entry point deliberately or accidentally left in a software program. It allows someone who knows about it to bypass the system's normal authentication and security protocols. This can be used by developers for troubleshooting but is a major security risk if discovered and exploited by malicious actors to gain access.

195. (A) Allows an administrator to switch to a different user's session without logging them out.
Fast User Switching is a Windows feature that enables multiple users to have active sessions on a single computer at the same time. It allows one user to switch to their account while another user's programs and files remain open and running in the background. This improves convenience in shared environments by avoiding the need for users to log out completely.

196. (B) Linker
A linker is a system software tool that plays a crucial role in the development process. It takes multiple object files, which are produced by a compiler or assembler, and combines them into a single executable program. During this process, the linker resolves all symbolic references between the files and integrates necessary libraries, ensuring that the final program is ready for execution.

197. (C) Resource allocation
Operating systems provide various services to ensure that the computer system operates efficiently. While user interfaces and program execution are focused on the user, resource allocation is a system-level service. It manages the distribution of CPU time, memory, and I/O devices among multiple processes, ensuring that resources are used effectively and that no single process can monopolize the entire system.

198. (B) 1, 3, and 4 only
Synchronization tools are essential for managing concurrent processes. A semaphore uses atomic operations to control access, and a mutex lock serves as a binary semaphore for mutual exclusion. To improve efficiency, blocking semaphores put waiting processes into a queue, avoiding the CPU waste of busy waiting. While useful, Peterson's solution is a software method, not a hardware-based one, for synchronization.

199. (B) THE question of whether it duplicates only THE calling thread or all threads.
In a multithreaded program, the behavior of the `fork()` system call is a significant design issue. The fundamental question is whether the new child process should contain a duplicate of every thread from the parent or only the single thread that made the call. Modern systems often provide different versions of the call or specific rules to handle this complexity.

200. (C) Dispatch latency
Dispatch latency is the time taken by the dispatcher to stop the currently running process and start another one on the CPU. This delay is an important factor in the responsiveness of a system, especially in real-time environments. It includes the time required to perform a context switch, which involves saving the state of the old process and loading the state.

201. (B) P-4, Q-1, R-3, S-2
Deadlock formation depends on four specific conditions. Mutual exclusion means resources are non-sharable. Hold and wait describes a process waiting for more resources while keeping others. No preemption prevents the OS from taking resources away. Circular wait involves a closed chain of processes each waiting for the next. Together, these conditions create a permanent system stall that requires specific handling or prevention.

202. (B) Logical address space
The logical address space is the set of all virtual addresses that a program can generate during its execution. This space is private to the process and is isolated from other processes. The operating system, with the help of the memory management unit, maps this logical space onto the physical address space of the actual RAM, providing flexibility and security for software.

203. (D) Mounted Access
File access methods define how data within a file is retrieved. Sequential access reads data in order, direct access allows jumping to any block, and indexed access uses a table for fast lookup. Mounting is not an access method; it is

Click Here to Buy book on Amazon

such as input and output operations, file management, or memory allocation tasks safely.

258. (A) The kernel.
The kernel represents the central core of an operating system, performing the most vital low-level functions required for system operation. It is responsible for critical tasks such as managing active processes, allocating system memory, handling file systems, and controlling input or output devices. As the primary component, it remains in memory throughout the entire duration of the system operation.

259. (C) Loading the operating system kernel into memory and initializing the system.
The system boot is a critical startup sequence managed by firmware that initiates when a computer is powered on. Its main responsibility is to locate and load the operating system kernel into the main memory. Once loaded, the kernel initializes all necessary hardware components, device drivers, and system services, preparing the environment for user interaction and application execution later.

260. (A) Ensure system integrity by providing controlled access to hardware.
System calls are implemented to provide a secure and controlled gateway for user programs to access sensitive hardware resources. By requiring programs to request services through the kernel, the operating system can verify permissions and prevent unauthorized or harmful actions. This structure is essential for maintaining overall system integrity and preventing individual applications from crashing the entire computer system.

261. (D) To allow cooperating processes to exchange data and information.
Interprocess communication provides the necessary mechanisms for separate processes to share data and synchronize their activities. In modern operating systems, cooperating processes often need to exchange information to perform complex tasks or coordinate their execution. Models like client-server systems utilize these communication protocols to facilitate efficient interactions between various software components running concurrently on the machine for users.

262. (B) Critical section
A critical section refers to a specific portion of a program's code where shared resources, such as variables or files, are accessed. To prevent data corruption or inconsistency, only one process should be allowed to execute within its critical section at any given time. Managing access to these code segments is a fundamental challenge in concurrent programming and synchronization tasks.

263. (C) Software-based solution to the critical-section problem for two processes.
Peterson's Solution is a classic algorithm developed to solve the critical-section problem through software-based logic. It specifically addresses the needs of two processes that must share resources without conflict. By using shared variables for turn-taking and indicating interest, it ensures mutual exclusion, progress, and bounded waiting, making it a foundational concept in the study of process synchronization today.

264. (D) By using a shared integer variable and atomic operations to control access to a resource.
Semaphores are powerful synchronization primitives that use an integer variable to control access to shared resources. They are accessed exclusively through two atomic operations, commonly known as wait and signal. By incrementing or decrementing the value, semaphores manage process entry into critical sections, ensuring that resource limits are respected and preventing multiple processes from causing data conflicts during execution.

265. (A) The concurrent access to shared resources by multiple processes.
The critical-section problem arises when multiple processes execute concurrently and attempt to manipulate shared data or hardware resources simultaneously. Without proper synchronization, these overlapping operations can lead to race conditions where the final result depends on the specific order of execution. Maintaining data consistency requires strictly controlled access to these shared areas of the program code during runtime.

266. (B) The process's code section and data section.

NODIA

acquire the resources it needs and finish its execution. However, the presence of a cycle only indicates a potential deadlock, depending on the number of instances.

231. (B) Demand Paging
Demand paging is an efficient virtual memory management technique where pages of a process are only loaded into physical memory when they are actually needed during execution. This reduces the initial loading time and saves memory, as parts of the program that are never used stay on the disk. A page fault is triggered when the process accesses a missing page.

232. (D) RAID 10 (1+0)
RAID 10 is a hybrid configuration that combines the benefits of both RAID 1 and RAID 0. It uses mirroring to provide redundancy and striping to enhance performance. By creating a stripe of mirrored disk sets, it offers high data reliability and fast read and write speeds. This makes it an excellent choice for critical, high-performance applications like large-scale database servers.

233. (A) A Virtual File System (VFS) layer.
The Virtual File System is an abstraction layer in the operating system kernel that provides a uniform interface for different file system types. It allows applications to use the same system calls, such as read and write, regardless of whether the underlying storage is an ext4, NTFS, or a network-based file system. This modularity significantly improves the system's flexibility and portability.

234. (A) Both Statement I and Statement II are correct.
The access matrix is a fundamental theoretical model for security, but directly implementing it as a large table is inefficient and impractical. Instead, modern systems use more compact implementations like access control lists, which store entries for each object, or capability lists, which store entries for each user or domain. These practical methods efficiently manage permissions while adhering to the core principles.

235. (B) Hide hardware differences from upper levels of THE system.
The Hardware Abstraction Layer is a critical component of the Windows operating system that provides a standardized interface between the physical hardware and the kernel. By hiding the specific details of different motherboard architectures and processor types, it allows the rest of the OS and its drivers to remain portable. This design ensures that Windows can run on diverse hardware platforms.

236. (C) Assembler
An assembler is a specialized system software tool that translates programs written in assembly language into machine code. It takes the human-readable mnemonics and symbolic labels of the assembly source and converts them into the binary instructions that the CPU can directly execute. This translation is a necessary step in the development of low-level software like device drivers and kernels.

237. (C) Selecting processes from THE job pool on disk to be loaded into memory.
The long-term scheduler, also known as the job scheduler, manages the degree of multiprogramming in a system. It selects processes from a pool of jobs stored on the disk and decides which ones should be brought into main memory to enter the ready queue. By controlling the mix of processes, it ensures that system resources like the CPU and memory are.

238. (D) 1, 2, 3, and 4
Operating system security involves understanding various threats and defense mechanisms. Standalone worms and host-dependent viruses are common program threats. Cryptography provides powerful tools for securing data and verifying user identities. To manage permissions, the access matrix model defines the specific rights each process has over various system objects, ensuring that resources are accessed only by authorized and authenticated users.

239. (D) All of THE above.
Signal handling in a multithreaded environment offers several delivery options. A signal can be sent to the specific thread it applies to, to every thread in the process, or only to certain designated threads. Modern operating systems often provide a combination of these methods to allow developers to manage notifications effectively while maintaining the stability and correct behavior of complex, concurrent programs.

Downloaded from www.pdf.tube

275. (C) Thread scheduling
In modern operating systems that support multithreading, the scheduler often focuses on kernel-level threads rather than treating entire processes as single units. This practice, known as thread scheduling, allows the CPU to switch between different execution paths within the same process or across different processes. This granular approach is essential for maximizing performance on modern multicore hardware architectures.

276. (A) Resource preemption
For a deadlock state to exist, four specific conditions must occur simultaneously: mutual exclusion, hold and wait, no preemption, and circular wait. Resource preemption is actually a mechanism used to recover from deadlocks, rather than a requirement for them to happen. If a system allows preemption, it can break the no preemption condition and successfully prevent a deadlock.

277. (D) Deadlock prevention
Deadlock prevention is a proactive strategy that involves setting system rules to ensure that at least one of the four necessary deadlock conditions can never occur. By negating conditions like hold and wait or circular wait, the system guarantees that a deadlock state is impossible. This often requires processes to request all resources at once or follow strict ordering rules.

278. (B) Prior information about the future resource needs of processes.
Deadlock avoidance is a dynamic method where the operating system evaluates each resource request to ensure the system remains in a safe state. To perform this assessment, the OS requires advance knowledge regarding the maximum resources each process might eventually need. With this information, the system can determine if granting a current request could potentially lead to a deadlock.

279. (C) To identify that a deadlock has already occurred in the system.
A deadlock detection algorithm is used in systems that allow processes to potentially enter a deadlocked state. Its role is to periodically examine the status of resource allocations and process requests to identify if a cycle of waiting has formed. Once a deadlock is detected, the system must then initiate a recovery procedure to resolve the resource conflict.

280. (D) A recovery scheme to break the existing deadlock.
After a deadlock detection algorithm identifies that a deadlock exists, the system must apply a recovery scheme to break the cycle. This typically involves terminating one or more of the deadlocked processes or preempting resources from them. The goal of this recovery process is to return the system to a normal state where all remaining processes can proceed safely.

281. (A) Paging uses fixed-sized blocks (pages/frames), whereas segmentation uses variable-sized logical units.
Paging and segmentation are two distinct memory management techniques. Paging divides physical and logical memory into uniform, fixed-size blocks called frames and pages. In contrast, segmentation divides memory based on the logical structure of a program, using variable-sized units like code segments, data segments, and stacks. This approach aligns more closely with how programmers view their code during development.

NODIA

282. (B) Demand paging
Demand paging is a virtual memory implementation where pages are loaded into the physical memory only when they are explicitly requested during execution. This technique optimizes memory usage by ensuring that only the currently active portions of a program occupy RAM. It reduces initial loading time and the amount of I/O required, although it may trigger occasional page faults.

283. (B) A high level of paging activity where the system spends more time on I/O than on execution.
Thrashing occurs when an operating system spends a significant majority of its time swapping pages in and out of memory rather than executing instructions. This state typically happens when processes do not have enough physical frames to hold their active working sets. The resulting high page fault rate causes a sharp drop in system performance and CPU utilization levels.

284. (A) Select a page in memory to be replaced when a new page must be loaded.
When a page fault occurs and no free memory frames are available, the operating system must use a page replacement algorithm. This algorithm is responsible for selecting an existing page in memory to be evicted, making room for the newly required page. Effective algorithms aim to choose pages that are unlikely to be needed again in the near future.

285. (C) By mapping a disk file directly into a process's virtual address space.
Memory-mapped files improve I/O performance by allowing a portion of a file on the disk to be mapped directly into the virtual address space of a process. This enables the program to treat file data as if it were part of its own memory, using standard memory access instructions instead of explicit read and write system calls for efficiency.

286. (D) To minimize seek time and rotational latency for I/O requests.
Disk scheduling algorithms are designed to optimize the performance of mass-storage devices by organizing the order in which I/O requests are serviced. The primary objective is to minimize the movement of the disk head, known as seek time, and the time spent waiting for the correct sector, known as rotational latency. Efficient scheduling significantly improves overall system throughput.

287. (C) Improve the performance and reliability of disk storage by combining multiple disks.
RAID technology combines multiple physical disk drives into a single logical unit to enhance data storage performance and provide redundancy. By distributing data across several disks, it can increase read and write speeds. Furthermore, the inclusion of redundant information allows the system to remain operational and recover data even if an individual physical disk drive fails during operation.

288. (B) Indexed allocation
The passage identifies several key methods for allocating disk blocks to files, including contiguous, linked, and indexed allocation. In an indexed allocation scheme, each file has its own index block that contains pointers to all its specific data blocks on the disk. This method supports efficient direct access to any part of the file without requiring contiguous space.

289. (D) File system mounting
Mounting is the essential process of making a file system on a storage device accessible to the operating system's overall directory structure. By attaching the file system to a specific mount point, the OS integrates the device's contents into the existing file tree. This allows users and applications to interact with the storage using standard file paths and navigation.

290. (A) Physical properties of the storage devices.
The file system serves as an abstraction layer that presents a simplified, logical view of data storage to users and applications. It hides the complex physical details of the storage hardware, such as the organization of tracks, sectors, and platters. This allows developers to work with files and directories without needing to manage the underlying hardware mechanics of disks.

291. (C) Protection is a mechanism to control resource access, while security is a broader defense against all attacks.
Protection focuses on the internal mechanisms that control how processes and users access specific system resources. Security is a broader concept encompassing the defense of the entire system against various internal and external threats, such as malware or network attacks. While protection manages permissions within the system, security ensures the system remains safe from malicious exploitation and unauthorized intrusion.

292. (A) The access rights of a subject over an object.
The Access Matrix is a theoretical model used to define protection policies in an operating system. In this matrix, the rows represent active subjects like users or processes, and the columns represent passive objects like files or hardware devices. Each entry in the matrix specifies the particular set of access rights that a subject possesses over a given object correctly.

293. (D) A denial-of-service attack
System and network threats involve attacks that target the infrastructure and availability of computing resources. A denial-of-service attack is a prime example, where an attacker attempts

Click Here to Buy book on Amazon

to make a service unavailable by overwhelming it with traffic or requests. Unlike program threats that infect individual files, network threats exploit communication protocols and system vulnerabilities to disrupt service for many users simultaneously.

294. (B) Encrypting data to ensure confidentiality and integrity.
Cryptography is a fundamental security tool used to protect sensitive information by transforming it into an unreadable format. Its primary purpose is to ensure data confidentiality, preventing unauthorized parties from understanding the content. Additionally, it helps maintain data integrity by providing mechanisms to detect if the information has been altered or tampered with during transmission or storage.

295. (C) To verify a user's identity before allowing access to the system.
User authentication is the critical process of verifying that a person attempting to access a system is indeed who they claim to be. This is typically achieved through credentials like passwords, biometrics, or security tokens. By establishing a user's identity before granting access, the operating system can ensure that protection policies and access rights are applied correctly.

296. (D) The ability to dynamically load and unload kernel modules.
Although Linux utilizes a monolithic kernel design where most services run in a single address space, it achieves significant flexibility through the use of kernel modules. These modules can be dynamically loaded or unloaded while the system is running. This allows for the addition of new device drivers or file systems without requiring a complete system reboot or kernel recompilation.

297. (A) Windows runs many OS services as system components, similar to a microkernel approach.
Windows employs an architectural design that is more closely aligned with a microkernel philosophy compared to the monolithic approach of Linux. In Windows, many operating system services run as separate system components in user mode rather than being integrated into the core kernel. This modular structure helps isolate different services, potentially improving system stability and providing support for specialized features.

298. (B) Robustness
Robustness is a key design goal for distributed systems, referring to the system's ability to continue operating correctly even when one or more of its individual nodes fail. Achieving robustness requires sophisticated mechanisms for fault detection and recovery. It ensures that the overall system remains available to users despite the inevitable hardware or software problems occurring in a large network.

299. (B) The NT File System (NTFS).
The NT File System is a sophisticated and advanced file system developed specifically for the Windows operating system family. It provides various features such as disk quotas, file compression, and security permissions that are not found in simpler file systems. NTFS is a core component of the Windows storage architecture, designed to handle large volumes and complex data management tasks.

300. (A) Managing communication structures and achieving fault tolerance.
Designing a distributed system is significantly more complex than creating a centralized one due to several unique challenges. Developers must manage communication structures and protocols across a network while ensuring the system remains robust against individual node failures. Achieving transparency and scalability in an environment with multiple autonomous computers requires intricate coordination and fault-tolerant architectural designs.

Downloaded from www.pdf.tube

CHAPTER 06

Software Engineering

1. Which of the following best defines a Software Process?
(A) A set of software tools and notations used to build software products.
(B) A collection of activities and tasks required for creating a quality product.
(C) A specific programming language and compiler chosen for a given project.
(D) The final executable code and documentation delivered to the customer.

2. Consider the following statements regarding software engineering terms:
Statement I : A Task Set defines the actual work to be done to accomplish the objectives of a software engineering action.
Statement II : A Framework Activity is a high-level stage in the software process, such as Communication or Deployment.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

3. Which of the following is NOT considered a generic framework activity in a software process model?
(A) Planning
(B) Modeling
(C) Marketing
(D) Construction

4. In the context of software engineering, what is a Process Pattern?
(A) A reusable graphical user interface design template for applications.
(B) A standardized algorithm for data encryption and secure storage.
(C) A legal document outlining the terms of software licensing and use.
(D) A template describing a recurring problem and a proven solution.

5. Identify the incorrect statement regarding Prescriptive Process Models.
(A) They strive for structure and order in software development.
(B) They prescribe a distinct set of process elements and a predictable process workflow.
(C) They are designed to be highly adaptive to rapid changes in requirements.
(D) The waterfall model is a classic example of a prescriptive model.

NODIA

6. Match the Agile Process Models in List I with their Key Characteristics in List II:
List I: (Agile Process Model)
P. Extreme Programming (XP)
Q. Scrum
R. Crystal
S. Feature Driven Development (FDD)
List II: (Key Characteristic)
1. Emphasizes developing a model of features and then a feature list.
2. Focuses on maneuverability and is treated as a family of methodologies.
3. Uses short, iterative cycles called sprints managed by a specific role.
4. Emphasizes values like communication, simplicity, and courage, with practices like pair programming.
Codes:
(A) P-4, Q-3, R-2, S-1
(B) P-3, Q-4, R-1, S-2
(C) P-4, Q-1, R-3, S-2
(D) P-2, Q-3, R-4, S-1

7. What is the primary focus of Aspect-Oriented Software Development (AOSD)?
(A) Developing software based on pre-built software components.
(B) Using mathematical models to prove the correctness of software.
(C) Designing and developing applications for the World Wide Web.

(D) Identifying, specifying, and representing cross-cutting concerns.

8. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The Component-Based Development model promotes software reuse.
Reason (R) : It emphasizes the construction of systems using existing software components that have been built for reuse.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

9. Consider the following activities associated with the generic process model:
I. Planning
II. Deployment
III. Construction
IV. Marketing
V. User Training
Which of the above statement(s) is/are correct?
(A) I, II, and V only
(B) II, III, and IV only
(C) I, II, and III only
(D) I, IV, and V only

10. Which process model follows a strict linear sequence of phases where the output of one phase becomes the input for the next?
(A) The Waterfall Model
(B) The Spiral Model
(C) The Prototyping Model
(D) The Scrum Model

11. Consider the following statements about Agile Process Models:
I. They emphasize incremental delivery of software.
II. They favor rigid planning and extensive documentation at the beginning of the project.
III. They adapt well to changing requirements throughout the development lifecycle.
IV. Customer collaboration is valued over contract negotiation.
Which of the above statements are correct?
(A) I, II, and III
(B) I, III, and IV
(C) II and IV
(D) I and II

12. In the context of Formal Methods in software engineering, what is the main objective?
(A) To conduct extensive user interviews to gather requirements.
(B) To use mathematical specifications for designing and verifying a system.
(C) To create a detailed project timeline using Gantt charts.
(D) To automate the generation of user interface code.

13. Which agile methodology uses a fixed-length iteration, called a sprint, to create a potentially shippable product increment?
(A) Extreme Programming (XP)
(B) Crystal
(C) Scrum
(D) Feature Driven Development (FDD)

14. The Dynamic System Development Model (DSDM) is an agile model that is built around which core principle?
(A) A strict adherence to the initial project plan without any deviation.
(B) The exclusion of the user from the development process to speed up delivery.
(C) The use of formal mathematical proofs for every module.
(D) The 80/20 rule, where 80% of the application can be delivered in 20% of the time.

15. What is the primary role of a Scrum Master in the Scrum framework?
(A) To facilitate the Scrum process and remove impediments for the development team.
(B) To dictate the technical design of the software.
(C) To manage the project budget and timeline.
(D) To write and execute all test cases for the sprint.

16. Which of the following is NOT a core value of Extreme Programming (XP)?
(A) Simplicity
(B) Comprehensive Documentation
(C) Communication

Click Here to Buy book on Amazon

(D) Feedback

17. Arrange the following phases of a generic waterfall model in chronological order:
I. System/Information Engineering and Modeling
II. Requirements Analysis
III. Design
IV. Code Generation
V. Testing
(A) I, II, III, IV, V
(B) II, I, III, IV, V
(C) I, III, II, V, IV
(D) II, III, I, IV, V

18. What does Web Engineering specifically address?
(A) Systematic and disciplined approaches to developing high-quality Web systems.
(B) The engineering of general-purpose desktop applications for enterprise users.
(C) The process of designing and manufacturing specialized computer hardware.
(D) The study of network protocols, infrastructure, and data routing.

19. Identify the correct pair from the following options related to software process models.
(A) Scrum - A prescriptive model emphasizing documentation.
(B) Waterfall - An agile model focused on iterative development.
(C) Adaptive Software Development (ASD) - An agile model focused on collaboration and learning cycles.
(D) Formal Methods - A testing methodology using black-box techniques.

20. A hospital needs a new patient management system. The requirement The system must allow a doctor to access a patient's medical history is an example of what type of requirement?
(A) Non-Functional Requirement
(B) Design Constraint
(C) Performance Requirement
(D) Functional Requirement

21. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Non-functional requirements define the quality attributes, performance, and security of a system.
Reason (R) : They describe how a system should perform a certain function, rather than what the system should do.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

22. Which technique for eliciting requirements involves creating scenarios of how a user will interact with the system to achieve a specific goal?
(A) Formal specification
(B) Developing Use Cases
(C) Quality Function Deployment
(D) Stakeholder interviews

23. Consider the following statements about a Software Requirements Specification (SRS) document:
I. It serves as a contract between the client and the developer.
II. It only includes functional requirements.
III. It is created after the software has been coded and tested.
IV. It should be unambiguous, consistent, and complete.
V. It is primarily used by the marketing team for promotional activities.
Which combination of the above statements is correct?
(A) I and IV
(B) II and III
(C) I, III, and V
(D) II, IV, and V

24. The requirement The system should be able to handle 1000 concurrent users with a response time of less than 2 seconds is an example of a:
(A) Non-Functional Requirement
(B) Functional Requirement
(C) Design Requirement
(D) Structural Requirement

25. What is the primary purpose of a Requirements Review?
(A) To start coding the software immediately.

(B) To finalize the project budget and schedule.
(C) To assess the quality and completeness of the requirements specification.
(D) To market the product to potential customers.

26. Match the Requirement Elicitation Techniques in List I with their Descriptions in List II:
List I: (Requirement Elicitation Technique)
P. Interviews
Q. Use Cases
R. Prototypes
S. Questionnaires
List II: (Description)
1. Building a working model of the system to get feedback.
2. A set of written questions distributed to a large number of stakeholders.
3. Direct discussion with stakeholders to gather information.
4. Describing system behavior from a user's perspective.
Codes:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-3, R-1, S-4

27. Which of the following is NOT a characteristic of a good Software Requirements Specification (SRS)?
(A) Verifiable
(B) Modifiable
(C) Traceable
(D) Ambiguous

28. The process of discovering, analyzing, documenting, and validating the requirements for a software system is known as:
(A) Software Design
(B) Requirements Engineering
(C) Software Testing
(D) Software Maintenance

29. Consider the following statements regarding software requirements:
Statement I : Functional requirements describe the services the software should provide.
Statement II : Non-functional requirements are constraints on the services or functions offered by the system.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

30. In the context of software design, what does Abstraction refer to?
(A) The degree to which the elements inside a module belong together.
(B) The measure of interdependence between software modules.
(C) The process of breaking down a system into its constituent parts.
(D) Hiding implementation details of a module to show functionality.

31. A design pattern in software engineering is best described as:
(A) A reusable solution to a common problem within a specific design context.
(B) A finished piece of code that can be directly used in a program.
(C) A formal mathematical proof of a program's correctness.
(D) A diagram showing the physical deployment of a system.

32. The design principle that suggests a system should be decomposed into parts with minimal overlap in functionality is known as:
(A) Information Hiding
(B) Separation of Concerns
(C) Functional Independence
(D) Architectural Pattern

33. Identify the key principles of effective software design from the following list:
I. Modularity
II. High Cohesion
III. Low Coupling
IV. Information Hiding
Which of the above are correct?
(A) I and II only
(B) II, III, and IV only
(C) I, II, and III only
(D) I, II, III, and IV

34. What is the relationship between Cohesion and Coupling in good software design?
(A) Good design aims for high cohesion and low coupling.

Downloaded from www.pdf.tube

(A) Correctness
(B) Portability
(C) Reusability
(D) Integrity

53. A project team decides not to include a complex, high-risk feature in the initial release of a software product. This is an example of which risk management strategy?
(A) Risk Transfer
(B) Risk Acceptance
(C) Risk Avoidance
(D) Risk Monitoring

54. The LOC (Lines of Code) metric is a common method for:
(A) Measuring software quality
(B) Software Sizing
(C) Determining user satisfaction
(D) Testing software reliability

55. Arrange the typical steps of project risk management in a logical sequence:
I. Risk Identification
II. Risk Analysis
III. Risk Mitigation Planning
IV. Risk Monitoring
(A) I, II, III, IV
(B) II, I, IV, III
(C) I, III, II, IV
(D) III, I, II, IV

56. Function Point (FP) analysis is a method for software sizing that is based on:
(A) The number of lines of source code.
(B) The complexity of the algorithm used.
(C) The functionality of the software from the user's point of view.
(D) The number of developers on the team.

57. The COCOMO model is used for:
(A) Software testing
(B) Software design
(C) Estimating project cost and effort
(D) Managing software configuration

58. Consider the following list of project metrics:
I. Number of external inputs
II. Number of external outputs
III. Number of logical internal files
IV. Number of developers
V. Number of external interface files
Which of the above are considered valid inputs for Function Point (FP) calculation?
(A) I, II, and IV only
(B) I, II, III, and V only
(C) III, IV, and V only
(D) I, IV, and V only

NODIA

59. What is a primary disadvantage of using Lines of Code (LOC) as a software sizing metric?
(A) It is easy to calculate.
(B) It is language-dependent and can be misleading for measurement.
(C) It can be measured early in the development lifecycle.
(D) It is directly related to the functionality of the software.

60. A project manager is creating a visual representation of project tasks along a horizontal timeline. This is most likely a:
(A) Use Case Diagram
(B) Data Flow Diagram
(C) Time-line Chart (Gantt Chart)
(D) Class Diagram

61. The Constructive Cost Model (COCOMO) has different modes based on project complexity. Which of the following is NOT a COCOMO mode?
(A) Organic
(B) Semi-detached
(C) Embedded
(D) Fragmented

62. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Function Point (FP) based estimation is generally considered more reliable than LOC based estimation.
Reason (R) : FP is language-independent and is based on the requirements of the software, which are known earlier in the lifecycle.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

44. Consider the following statements regarding Software Quality:
I. Quality Control (QC) is a process that focuses on identifying defects in the actual products.
II. Quality Assurance (QA) is a set of activities for ensuring quality in the processes by which products are developed.
III. QA is a subset of QC.
IV. Testing is a primary activity of Quality Control.
Which of the above statements are correct?
(A) I, II, and IV
(B) I, II, and III
(C) II and III only
(D) I and IV only

45. The ISO 9126 quality standard defines six main quality characteristics. Which of the following is NOT one of them?
(A) Functionality
(B) Reliability
(C) Usability
(D) Affordability

46. What is the primary focus of Risk Management in a software project?
(A) To eliminate all possible risks before the project starts.
(B) To identify, analyze, and control potential problems impacting the project.
(C) To transfer all project risks to the client.
(D) To focus solely on technical risks and ignore business risks.

47. A risk mitigation strategy that involves developing a contingency plan in case a foreseen risk becomes a reality is known as:
(A) Risk Avoidance
(B) Risk Transfer
(C) Risk Monitoring
(D) Risk Contingency Planning

48. Match the ISO 9126 Quality Factors in List I with their Descriptions in List II:
List I: (ISO 9126 Quality Factor)
P. Portability
Q. Efficiency
R. Maintainability
S. Reliability
List II: (Description)
1. The capability of the software to provide its functions under stated conditions for a specified period of time.
2. The ease with which a software system can be modified to correct faults or improve performance.
3. The performance of the software relative to the amount of resources used.
4. The ability of software to be transferred from one environment to another.
Codes:
(A) P-4, Q-3, R-2, S-1
(B) P-3, Q-4, R-1, S-2
(C) P-4, Q-1, R-3, S-2
(D) P-1, Q-2, R-4, S-3

49. What is Software Reliability defined as?
(A) The speed at which the software executes its functions.
(B) The ease with which the software can be used by a novice.
(C) The number of lines of code in the software system.
(D) The probability of failure-free operation in an environment.

50. The RMMM plan in risk management stands for:
(A) Risk Mitigation, Monitoring, and Management
(B) Risk Measurement, Mitigation, and Management
(C) Resource Mitigation, Monitoring, and Management
(D) Risk Mitigation, Measurement, and Monitoring

51. Consider the following statements regarding software quality:
Statement I : Quality Assurance (QA) is a proactive process aimed at preventing defects.
Statement II : Quality Control (QC) is a reactive process aimed at finding defects.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

52. Which of McCall's quality factors pertains to the effort required to transfer a program from one hardware or software environment to another?

Click Here to Buy book on Amazon

(B) Good design aims for high cohesion and high coupling.
(C) Good design aims for low cohesion and low coupling.
(D) Good design aims for low cohesion and high coupling.

35. A module that performs a single, well-defined task or function is an example of which type of cohesion?
(A) Coincidental Cohesion
(B) Logical Cohesion
(C) Functional Cohesion
(D) Temporal Cohesion

36. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Low coupling is a desirable characteristic in software design.
Reason (R) : Low coupling reduces the ripple effect, where a change in one module causes errors or requires changes in other modules.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

37. What is the primary purpose of Architectural Design in the software development lifecycle?
(A) To define the detailed logic for each individual component.
(B) To design the database schema and data structures.
(C) To describe the system structure, components, and relationships.
(D) To create the look and feel of the user interface.

38. The principle of Information Hiding suggests that:
(A) All information within a software project should be encrypted.
(B) Modules should be designed so that internal information is inaccessible to those without need.
(C) Users should not be shown error messages.
(D) Source code should not be commented to protect intellectual property.

39. Which of the following best describes Component Level Design?
(A) It defines the overall structure of the software architecture.
(B) It focuses on the aesthetics and usability of the front end.
(C) It designs the data structures required by the software.
(D) It transforms structural elements into procedural descriptions.

40. Identify the incorrect pair regarding software design concepts.
(A) Modularity - The process of hiding the internal details of a module.
(B) Cohesion - A measure of the degree to which elements of a module are functionally related.
(C) Coupling - A measure of the degree of interdependence between modules.
(D) Abstraction - Focusing on essential characteristics without irrelevant details.

NODIA

41. A banking application has separate modules for 'customer management', 'loan processing', and 'account transactions'. This decomposition is an example of applying which design principle?
(A) Information Hiding
(B) Cohesion
(C) Coupling
(D) Modularity

42. Which of the following is a key goal of User Interface (UI) Design?
(A) To choose the most complex algorithms for backend processing.
(B) To make user interaction as simple and efficient as possible.
(C) To ensure the database is normalized to the highest form.
(D) To write the most optimized machine code.

43. Which of McCall's Quality Factors refers to the effort required to learn, operate, and interpret the output of a program?
(A) Reliability
(B) Efficiency
(C) Usability
(D) Portability

63. The primary purpose of project scheduling in software project management is to:
(A) Select the programming language for the project.
(B) Allocate effort and assign tasks to specific team members.
(C) Perform unit testing on individual modules.
(D) Write the user manual for the software.

64. Which of the following estimation models is an example of an algorithmic cost estimation model?
(A) Expert judgment
(B) Analogy-based estimation
(C) COCOMO
(D) Parkinson's Law

65. Consider the following statements regarding project estimation:
Statement I : Effort estimation predicts the total amount of labor (e.g., in person-months) required to develop the software.
Statement II : Cost estimation predicts the financial expenditure required for the project, which is often derived from the effort estimation.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

66. What is the fundamental difference between Verification and Validation in software testing?
(A) Verification is Are we building the product right?, while Validation is Are we building the right product?.
(B) Verification is Are we building the right product?, while Validation is Are we building the product right?.
(C) Verification is done by the client, while Validation is done by the developer.
(D) Verification involves executing the code, while Validation involves reviewing documents.

67. An incorrect step, process, or data definition in a computer program is known as a(n):
(A) Failure
(B) Bug
(C) Error
(D) Fault

68. Consider the following statements about software testing:
I. A Bug is the observable incorrect behavior of the software.
II. A Fault is a static defect in the software (a mistake in the code).
III. An Error is a mistake made by a human (e.g., a developer).
IV. A Failure is the manifestation of a fault during execution.
Which of the above statements are correct?
(A) I and II only
(B) II, III, and IV only
(C) I, III, and IV only
(D) All of I, II, III, and IV

69. White-box testing is a testing technique where:
(A) The tester has no knowledge of the internal structure or code of the software.
(B) The tester has full knowledge of the internal structure and logic of the software.
(C) The testing is performed by the end-user at their own site.
(D) The focus is on the graphical user interface of the software.

70. Which testing technique is also known as glass-box or structural testing?
(A) Black-box testing
(B) White-box testing
(C) Alpha testing
(D) Beta testing

71. Match the Testing Terms in List I with their Descriptions in List II:
List I: (Testing Term)
P. Unit Testing
Q. Integration Testing
R. Alpha Testing
S. Beta Testing
List II: (Description)
1. Testing of individual software components or modules in isolation.
2. Testing performed by end-users at the developer's site.
3. Verifying the interfaces and interactions between integrated components.
4. Testing performed by end-users at their own site.
Codes:
(A) P-1, Q-2, R-3, S-4

NODIA

(A) Actions, tasks, and work products
(B) Programming languages
(C) Hardware configurations
(D) User manuals

132. The requirement The system shall use AES-256 encryption for all data transmissions is an example of a:
(A) Functional requirement
(B) Performance requirement
(C) Design constraint (a type of non-functional requirement)
(D) Usability requirement

133. Consider the following statements about software design:
I. Cohesion is a measure of the strength of functional relatedness of elements within a module.
II. Coupling is a measure of the interdependence among modules.
III. A good design should have low cohesion and high coupling.
IV. Information hiding conceals the implementation details of a module.
Which of the above statements are correct?
(A) I, II, and IV
(B) I, III, and IV
(C) II and III
(D) I and II only

134. According to ISO 9126, the ease with which a software system can be adapted to different specified environments is known as:
(A) Reliability
(B) Usability
(C) Portability
(D) Maintainability

135. What is the primary purpose of Staffing in the context of project scheduling?
(A) To estimate the total cost of the project.
(B) To assign people with the right skills to specific project tasks.
(C) To conduct performance reviews of the team members.
(D) To purchase hardware and software for the project.

136. Testing individual sub-programs, subroutines, or procedures in a software system is called:
(A) Integration Testing
(B) System Testing
(C) Unit Testing
(D) Acceptance Testing

137. Which of the following is a tool for version control in Software Configuration Management?
(A) Jira
(B) Git
(C) MS Project
(D) COCOMO

138. A developer team is unsure about the user interface requirements for a new system. Which process model would be most effective in this scenario to clarify requirements with the user?
(A) Waterfall Model
(B) Prototyping Model
(C) Formal Methods
(D) Cleanroom Software Engineering

139. In the context of use cases, who or what is an actor?
(A) Only a human user interacting with the system.
(B) The software developer writing the use case.
(C) Anyone or anything outside the system that interacts with the system.
(D) A specific module or component within the system.

140. Which design concept allows different objects to respond to the same message in different ways?
(A) Encapsulation
(B) Inheritance
(C) Polymorphism
(D) Abstraction

141. A quality assurance audit reveals that the development team is not following the coding standards defined in the project plan. This is a failure in:
(A) Quality Control
(B) Quality Assurance
(C) Risk Management
(D) Configuration Management

142. A project manager buys insurance to cover potential financial losses from a data breach. This is an example of which risk handling strategy?

Downloaded from www.pdf.tube

(D) Change Control

83. Analyzing a subject system to identify the system's components and their interrelationships and to create representations of the system in another form or at a higher level of abstraction is the definition of:
(A) Reverse Engineering
(B) Software Reuse
(C) Version Control
(D) Forward Engineering

84. Which of the following best describes Software Reuse?
(A) The process of fixing bugs in existing software.
(B) Building software systems using existing assets rather than from scratch.
(C) The process of discarding old software.
(D) The process of managing changes to software using version control.

85. Consider the following activities related to SCM:
I. Identifying configuration items
II. Establishing a baseline
III. Managing changes to the baseline
IV. Performing unit tests on code
V. Reporting the status of configuration items
Which of the above are considered core activities of Software Configuration Management?
(A) I, II, and IV only
(B) II, IV, and V only
(C) I, II, III, and V only
(D) I, III, and IV only

86. What is the main function of a Change Control Board (CCB)?
(A) To write the source code for new features.
(B) To approve or reject proposed changes to a software baseline.
(C) To test the software before release.
(D) To design the software architecture.

87. Which process model would be most appropriate for a project where requirements are well-understood and unlikely to change, and a formal, step-by-step approach is preferred?
(A) Scrum
(B) Extreme Programming (XP)
(C) Waterfall Model
(D) Adaptive Software Development

88. Identify the incorrect statement about Agile development.
(A) It promotes sustainable development pace.
(B) It values working software over comprehensive documentation.
(C) It requires a detailed, upfront design before any coding begins.
(D) It welcomes changing requirements, even late in development.

89. A developer team consists of a Product Owner, a Scrum Master, and the Development Team. This team structure is characteristic of which agile model?
(A) Crystal
(B) Scrum
(C) Feature Driven Development
(D) Extreme Programming

90. In Extreme Programming (XP), the practice of two programmers working together at one workstation is known as:
(A) Code Review
(B) Pair Programming
(C) Integration Testing
(D) Refactoring

91. Arrange the following software testing levels in the typical order they are performed:
I. Unit Testing
II. System Testing
III. Integration Testing
IV. Acceptance Testing
(A) I, III, II, IV
(B) I, II, III, IV
(C) IV, II, III, I
(D) I, III, IV, II

92. Which of the following statements accurately describes the relationship between an error, a fault, and a failure?
(A) A fault causes an error, which leads to a failure.
(B) A failure causes a fault, which leads to an error.
(C) A human mistake (error) leads to a defect (fault) causing a failure.
(D) A failure, a fault, and an error are all synonymous terms for a bug.

NODIA

112. The activity of checking if the software conforms to its specification is known as:
(A) Validation
(B) Verification
(C) Debugging
(D) Project Scheduling

113. Identify the correct statement regarding software design principles.
(A) Functional independence is achieved through low cohesion and high coupling.
(B) Information hiding makes modules more dependent on each other.
(C) Modularity facilitates parallel development and simplifies testing.
(D) Abstraction involves focusing on the detailed implementation of a component.

114. If a module's parts are grouped together because they are all executed in the same time span, what type of cohesion does it exhibit?
(A) Functional Cohesion
(B) Sequential Cohesion
(C) Temporal Cohesion
(D) Communicational Cohesion

115. A software team is developing a critical air traffic control system where safety and correctness are paramount. Which software engineering approach would be most appropriate?
(A) Extreme Programming (XP)
(B) Scrum
(C) Formal Methods
(D) Feature Driven Development

116. In the context of a project schedule, what is a milestone?
(A) A specific task assigned to a developer.
(B) A point in the timeline marking completion of a major deliverable.
(C) The total budget allocated for the project.
(D) A risk that has been identified but not mitigated.

117. Arrange the following COCOMO models in increasing order of project size and complexity:
I. Embedded
II. Semi-detached
III. Organic
(A) I, II, III
(B) III, II, I
(C) II, I, III
(D) III, I, II

118. Which of the following is NOT a type of black-box testing?
(A) Equivalence Partitioning
(B) Boundary Value Analysis
(C) Decision Table Testing
(D) Statement Coverage

119. A team analyzes a legacy system's source code to understand its functionality and architecture to document it for future maintenance. This process is an example of:
(A) Software Re-engineering
(B) Reverse Engineering
(C) Software Reuse
(D) Change Control

120. Match the Process Models in List I with their Primary Application Areas in List II:
List I: (Process Model)
P. Waterfall Model
Q. Prototyping Model
R. Spiral Model
S. Agile Scrum Model
List II: (Primary Application Area)
1. Large, complex, high-risk systems.
2. Projects where requirements are unstable or must be clarified.
3. Projects with well-defined, stable requirements.
4. Projects requiring rapid and incremental delivery with changing requirements.
Codes:
(A) P-1, Q-3, R-2, S-4
(B) P-3, Q-2, R-4, S-1
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-3, S-2

121. The term cross-cutting concerns such as logging, security, and transaction management, are primarily addressed by which development paradigm?
(A) Object-Oriented Design
(B) Component-Based Development
(C) Aspect-Oriented Software Development
(D) Structured Programming

122. Consider the following statements regarding software development artifacts:

Click Here to Buy book on Amazon

103. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).

Assertion (A) : Version control is a critical component of Software Configuration Management (SCM).

Reason (R) : It provides a mechanism to manage different versions of configuration items, allowing for tracking changes and reverting to previous states if necessary.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

104. Consider the following statements about software sizing metrics:

I. Lines of Code (LOC) is independent of the programming language used.
II. Function Points (FP) measure functionality delivered by the software.
III. LOC can be accurately measured at the requirements analysis stage.
IV. FP analysis is more subjective than counting LOC.

Which combination of the above statements is correct?

(A) I and III
(B) II and IV
(C) II only
(D) I, II, and IV

105. A measure of the degree of interdependence between two modules is the definition of:

(A) Cohesion
(B) Coupling
(C) Modularity
(D) Abstraction

106. In the Scrum framework, what is the Product Backlog?

(A) A list of all technical tasks to be completed in a sprint.
(B) An ordered list of everything that is known to be needed in the product.
(C) A list of bugs found during the last sprint.
(D) A chart showing the progress of the project.

107. The Crystal family of agile methodologies suggests that the specific practices of a project should be tuned based on:

(A) The programming language being used.
(B) The project's size and criticality.
(C) The geographical location of the development team.
(D) The number of lines of code.

108. Which of the following is a primary benefit of using a Component-Based Development model?

(A) Increased development time
(B) Reduced potential for software reuse
(C) Faster delivery and increased productivity
(D) Increased cost due to custom development

NODIA

109. What is a key characteristic of the Speculate, Collaborate, Learn cycle in Adaptive Software Development (ASD)?

(A) A rigid, unchangeable plan is created during the Speculate phase.
(B) Collaboration is limited to the development team only.
(C) Learning from mistakes and successes is a fundamental part of the process.
(D) The cycle is executed only once for the entire project.

110. Which process model is particularly well-suited for large, complex, and high-risk projects, as it incorporates risk analysis in each iteration?

(A) Waterfall Model
(B) Prototyping Model
(C) Scrum Model
(D) Spiral Model

111. Consider the following list of software development activities:

I. Develop an overall model
II. Build a feature list
III. Plan by feature
IV. Design by feature
V. Build by feature

Which of the above are characteristic steps of Feature Driven Development (FDD)?

(A) I and II only
(B) I, II, and III only
(C) III, IV, and V only
(D) I, II, III, IV, and V

(B) P-1, Q-3, R-2, S-4
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-3, R-1, S-4

72. Basis Path Testing is a technique of:
(A) Black-box testing
(B) White-box testing
(C) Performance testing
(D) Regression testing

73. The primary goal of Regression Testing is to:
(A) Test the software under heavy load.
(B) Test new features added to the software.
(C) Ensure changes have not introduced new faults in existing functionality.
(D) Test the software for the first time.

74. Black-box testing focuses on which aspect of the software?
(A) The internal logic and code paths.
(B) The functional requirements and specifications.
(C) The data structures used in the program.
(D) The efficiency of the algorithms.

75. What type of testing evaluates how a system performs under a particular workload, measuring attributes like response time and throughput?
(A) Stress Testing
(B) Regression Testing
(C) Performance Testing
(D) Unit Testing

76. Testing a system beyond its specified operational limits to observe its behavior is called:
(A) Performance Testing
(B) Stress Testing
(C) Alpha Testing
(D) Regression Testing

77. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Integration testing is performed after unit testing.
Reason (R) : Integration testing aims to find defects that arise from the interaction between different software modules that have already been tested individually.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

78. Control Structure Testing is a category of white-box testing that includes techniques like:
(A) Equivalence Partitioning
(B) Boundary Value Analysis
(C) Condition Coverage and Loop Testing
(D) Use Case Testing

79. Consider the following statements regarding software testing:
Statement I : Alpha testing is a form of acceptance testing performed in-house by the development or QA team.
Statement II : Beta testing involves releasing the software to a limited number of external users to obtain feedback.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

80. What is the main purpose of Software Configuration Management (SCM)?
(A) To manage and control changes to software artifacts throughout the lifecycle.
(B) To estimate the cost and effort of a software project.
(C) To design the user interface of the software.
(D) To test the performance and reliability of the software.

81. Which of the following is a primary activity of Software Configuration Management?
(A) Basis Path Testing
(B) Version Control
(C) Function Point Analysis
(D) User Interface Design

82. The process of creating a new version of a program from an existing one, by adding new functionalities or changing existing ones, is known as:
(A) Reverse Engineering
(B) Software Re-engineering
(C) Software Reuse

Downloaded from www.pdf.tube

Statement I : The SRS document is an internal document for the development team and is not shared with the customer.

Statement II : A use case describes a sequence of interactions between a user and a system to achieve a goal.

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

123. Consider the following list of requirements attributes:

I. Correctness
II. Ambiguity
III. Completeness
IV. Verifiability
V. Volatility

Which of the above characteristics are essential for a good Software Requirements Specification (SRS) document?

(A) I, III, and IV only
(B) II, IV, and V only
(C) I, II, and III only
(D) I, III, IV, and V only

124. What is the primary focus of Data Design in the software design phase?

(A) Creating the user interface layout.
(B) Defining the overall software architecture.
(C) Designing the data structures and database schema to be used.
(D) Writing the procedural logic for software components.

125. Which of the following is a key activity in Quality Assurance (QA) but not typically in Quality Control (QC)?

(A) Executing test cases.
(B) Finding bugs in the code.
(C) Defining the software development process and standards.
(D) Reviewing test results.

126. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).

Assertion (A) : Risk management is a continuous process throughout the project lifecycle.

Reason (R) : New risks can emerge, and the status of existing risks can change as the project progresses.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

127. If an estimation model is based on the formula Effort = a * (KLOC)^b, it is most likely a variant of which model?

(A) Function Point Model
(B) COCOMO Model
(C) Putnam Model
(D) Expert Judgment Model

128. What is the key difference between alpha and beta testing?

(A) Alpha testing uses white-box techniques, while beta testing uses black-box techniques.
(B) Alpha testing is done by developers, while beta testing is done by independent testers.
(C) Alpha testing is done at the developer's site, while beta testing is done at the customer's site.
(D) Alpha testing checks for functional requirements, while beta testing checks for non-functional requirements.

129. The process of modifying a software system to improve its maintainability without changing its functionality is known as:

(A) Reverse Engineering
(B) Re-engineering
(C) Forward Engineering
(D) Debugging

NODIA

130. Which of the following best describes Functional Independence in software design?

(A) The software can run on any operating system.
(B) The software is developed by a single programmer working independently.
(C) The software is designed with modules having high cohesion and low coupling.
(D) The software functions correctly without any user input.

131. A framework activity in a generic process model is a collection of related:

(D) Condition Coverage

152. The process of managing a 'baseline' and controlling changes to it is a central part of:
(A) Software Testing
(B) Requirements Elicitation
(C) Software Configuration Management
(D) Project Estimation

153. In the context of the Scrum agile model, what is a Sprint Retrospective?
(A) A meeting used to plan the technical work for the next sprint.
(B) A daily meeting for the team to synchronize activities.
(C) A meeting to inspect performance and create a plan for improvement.
(D) A demonstration of the completed work to the stakeholders.

154. What is the primary drawback of the Waterfall Model of software development?
(A) It promotes too much user involvement, slowing down the project.
(B) It is too flexible and lacks structure.
(C) It is rigid and does not easily accommodate requirement changes.
(D) It is only suitable for very small projects.

155. Which of the following principles is NOT part of the Agile Manifesto?
(A) Individuals and interactions over processes and tools
(B) Working software over comprehensive documentation
(C) Following a plan over responding to change
(D) Customer collaboration over contract negotiation

156. A software design specifies that a 'Customer' module directly accesses and modifies data within an 'Order' module. This is an example of:
(A) High cohesion
(B) Low cohesion
(C) High coupling
(D) Low coupling

157. Consider the following goals for User Interface Design:
I. To make the interface aesthetically pleasing.
II. To make the interface efficient and easy to use.
III. To ensure the interface works on all possible devices.
IV. To choose the backend database for the system.
Which of the above are correct goals of UI design?
(A) I and II only
(B) III and IV only
(C) I, II, and III only
(D) II and IV only

158. Consider the following factors:
I. Portability
II. Reliability
III. Cost
IV. Usability
V. Development Time
Which of the above factors are considered in McCall's Quality Model?
(A) I, II, and IV only
(B) I, II, III, and IV only
(C) III and V only
(D) I, II, and V only

159. A software development project has a small, highly skilled team and the requirements are not well-defined. This project fits the description for which COCOMO mode?
(A) Organic
(B) Semi-detached
(C) Embedded
(D) Detailed

160. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Regression testing is an essential part of software maintenance.
Reason (R) : Whenever a change is made to the software (for bug fixing or enhancement), regression testing is performed to ensure that the change has not adversely affected other parts of the software.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

Click Here to Buy book on Amazon

(A) Risk Avoidance
(B) Risk Mitigation
(C) Risk Transfer
(D) Risk Acceptance

143. Consider the following statements regarding software testing techniques:
Statement I : White-box testing requires the tester to have detailed knowledge of the program's internal logic.
Statement II : Black-box testing is performed without any knowledge of the internal workings of the software.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

144. Re-creating design and specification documents from existing source code is a primary goal of:
(A) Software reuse
(B) Change control
(C) Reverse engineering
(D) Project scheduling

145. What is the relationship between a Task Set and a Framework Activity?
(A) A Task Set is a more abstract concept than a Framework Activity.
(B) An activity is comprised of actions, and each action is defined by a Task Set.
(C) A Task Set and a Framework Activity are identical concepts.
(D) A Task Set is used for project scheduling, while a Framework Activity is used for cost estimation.

146. Which of the following is an example of a Functional Requirement?
(A) The system must be available 99.9% of the time.
(B) The user interface should be colored in shades of blue.
(C) The system shall allow a user to add items to a shopping cart.
(D) The system's response time must be less than 1 second.

147. A module where all processing elements are concentrated on one area of a data structure is exhibiting which type of cohesion?
(A) Functional Cohesion
(B) Communicational Cohesion
(C) Coincidental Cohesion
(D) Logical Cohesion

148. The 3 Cs of a user story in agile development are Card, Conversation, and Confirmation. What does Conversation refer to?
(A) The formal documentation of the requirement.
(B) The discussion between team and owner to clarify user story details.
(C) The negotiation of the project contract.
(D) The comments written in the source code.

149. According to McCall's quality model, Correctness and Reliability are factors related to:
(A) Product Revision
(B) Product Transition
(C) Product Operation
(D) Product Design

150. Match the Risk Terms in List I with their Examples in List II:
List I: (Risk Term)
P. Project Risk
Q. Product Risk
R. Business Risk
S. Technical Risk
List II: (Example)
1. The chosen database technology cannot handle the required transaction volume.
2. The developed software does not meet the quality standards and fails frequently.
3. The project team loses key personnel, affecting the schedule.
4. A competitor releases a similar product earlier, capturing the market.
Codes:
(A) P-3, Q-2, R-4, S-1
(B) P-1, Q-2, R-3, S-4
(C) P-3, Q-4, R-2, S-1
(D) P-2, Q-3, R-1, S-4

151. If a tester wants to ensure that every possible branch (true/false) from each decision point in the code is executed at least once, which testing coverage criterion are they targeting?
(A) Statement Coverage
(B) Branch/Decision Coverage
(C) Path Coverage

Downloaded from www.pdf.tube

93. A software project is characterized by a small, experienced team, rapidly changing requirements, and a need for frequent delivery of working software. Which process model would be the most suitable?
(A) Formal Methods Model
(B) Waterfall Model
(C) Agile Scrum Model
(D) Component-Based Development Model

94. What is the primary output of the Developing Use Cases activity in requirements engineering?
(A) A detailed project schedule.
(B) A description of interactions between an actor and the system.
(C) The final executable file.
(D) A database schema.

95. In Object-Oriented Design, bundling data (attributes) and methods (functions) that operate on the data into a single unit (a class) is known as:
(A) Inheritance
(B) Polymorphism
(C) Encapsulation
(D) Abstraction

96. Which McCall's quality factor relates to the extent to which a program can be reused in other applications?
(A) Portability
(B) Reliability
(C) Reusability
(D) Efficiency

97. What is the purpose of a baseline in Software Configuration Management?
(A) It is the initial version of the source code that is never changed.
(B) It is a formally accepted configuration version serving as a reference.
(C) It is a list of all bugs found during testing.
(D) It is the final version of the software delivered to the customer.

98. A team is tasked with improving the internal structure of existing code without changing its external behavior. This activity is called:
(A) Re-engineering
(B) Refactoring
(C) Reverse Engineering
(D) Regression Testing

99. Which of the following is an example of a Risk Monitoring activity?
(A) Identifying a list of potential risks at the start of a project.
(B) Assessing the probability and impact of each identified risk.
(C) Creating a plan to avoid a specific risk.
(D) Tracking risk indicators and triggers to see if a risk is becoming likely.

100. Match the Design Principles in List I with their Goals in List II:
List I: (Design Principle)
P. Separation of Concerns
Q. Functional Independence
R. Information Hiding
S. Modularity
List II: (Goal)
1. Decomposing a system into manageable and self-contained units.
2. Hiding the implementation details of a module from its users.
3. Breaking a complex problem into smaller, independent sub-problems.
4. Creating modules with high cohesion and low coupling.
Codes:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-2, S-1
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-3, R-1, S-4

NODIA

101. A requirement stating The user interface should be intuitive and easy to use is an example of which type of requirement?
(A) Functional
(B) Non-Functional (Usability)
(C) Architectural
(D) Data

102. What is the primary input for the Basis Path Testing technique?
(A) The software requirements specification.
(B) A flow graph of the program's control structure.
(C) A list of user stories.
(D) The project's risk management plan.

161. The process of creating an executable software system by analyzing, designing, and coding is referred to as:
(A) Reverse Engineering
(B) Forward Engineering
(C) Software Re-engineering
(D) Maintenance

162. A Process Lifecycle in software engineering refers to:
(A) The progression of a software process through its various stages.
(B) The warranty period of a software product.
(C) The time taken to compile the source code.
(D) The lifespan of the hardware on which the software runs.

163. The primary goal of Requirement Analysis and Modelling is to:
(A) Write the code for the system.
(B) Create models that refine and clarify the elicited requirements.
(C) Estimate the total cost of the project.
(D) Test the final product before delivery.

164. What does the term Separation of Concerns imply in software design?
(A) The development team and the testing team should work in separate buildings.
(B) A problem should be decomposed into distinct features with minimal overlap.
(C) User requirements should be separated from technical requirements.
(D) The database should be on a separate server from the application logic.

165. Consider the following principles of User Interface Design:
I. Place the user in control.
II. Reduce the user's memory load.
III. Make the interface consistent.
IV. Use complex and technical jargon.
Which of the above statements are correct?
(A) I and IV only
(B) II and III only
(C) I, II, and III only
(D) I, II, III, and IV

166. A risk is identified with a high probability of occurrence and a high impact on the project. What is the most likely action to be taken?
(A) Accept the risk and do nothing.
(B) Implement a mitigation or avoidance strategy immediately.
(C) Ignore the risk as it is too high to manage.
(D) Transfer the risk to a junior team member.

167. Arrange the following activities in the logical order of a software re-engineering process:
I. Code Restructuring
II. Reverse Engineering
III. Forward Engineering
IV. Data Restructuring
(A) II, I, IV, III
(B) I, II, III, IV
(C) II, IV, I, III
(D) III, I, IV, II

168. If a new version of software is released to a select group of external users for testing under real-world conditions, what is this process called?
(A) Alpha Testing
(B) Unit Testing
(C) Beta Testing
(D) Regression Testing

169. What is the most significant advantage of Function Point (FP) analysis over Lines of Code (LOC) for estimation?
(A) It is language-independent and technology-independent.
(B) It is quicker and easier to calculate.
(C) It can only be calculated after the code is written.
(D) It does not require knowledge of the system requirements.

170. In Object-Oriented Design, which concept enables a class to inherit properties and behavior from another class?
(A) Inheritance
(B) Abstraction
(C) Polymorphism
(D) Encapsulation

171. The Daily Scrum or Daily Stand-up is a short meeting for the Development Team to:
(A) Demonstrate the work completed to the Product Owner.
(B) Plan the entire project from start to finish.
(C) Synchronize activities and create a plan for the next 24 hours.

Click Here to Buy book on Amazon

interface.

193. The purpose of deriving test cases from a control flow graph in Basis Path Testing is to:
(A) Ensure every statement in the program is executed at least once.
(B) Guarantee every linearly independent path is executed at least once.
(C) Verify the non-functional requirements of the system.
(D) Test the user interface for usability issues.

194. Consider the following techniques for requirements gathering:
I. Interviews
II. Code reviews
III. Use cases
IV. Prototyping
V. Unit testing
Which of the above are considered valid techniques for requirements elicitation?
(A) I, II, and III only
(B) I, III, and IV only
(C) II, IV, and V only
(D) I, III, and V only

195. What does the separation of concerns principle in software design promote?
(A) Combining unrelated functionalities into a single module.
(B) Decomposing a large problem into smaller independent parts.
(C) Hiding implementation details from other parts of the system.
(D) Ensuring high dependency between different software modules.

196. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Agile models are iterative and incremental.
Reason (R) : They are designed to deliver software in small, functional pieces, allowing for flexibility and adaptation to change throughout the project lifecycle.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

197. Which of the following is a primary goal of Software Quality Assurance (SQA)?
(A) To find and fix as many bugs as possible before release.
(B) Ensuring development conforms to standards to prevent defects.
(C) To manage the project schedule and budget.
(D) To write the source code for the application.

198. The COCOMO model is an acronym for:
(A) Component Cost Model
(B) Constructive Cost Model
(C) Comprehensive Cost Model
(D) Constructive Component Model

199. What is the fundamental difference between an error and a fault in the context of software testing?
(A) They are synonymous terms.
(B) An error is a human mistake, while a fault is the defect in the code.
(C) A fault is a human mistake, while an error is the defect in the code.
(D) An error is an incorrect output, while a fault is a bug found by a user.

200. Match the SCM Activities in List I with their Descriptions in List II:
List I: (SCM Activity)
P. Version Control
Q. Change Control
R. Configuration Audit
S. Status Reporting
List II: (Description)
1. The process of managing and evaluating proposed changes to a baseline.
2. The process of tracking and managing multiple versions of software components.
3. Informing stakeholders about the current state of configuration items and changes.
4. Verifying that the software product matches the configuration items in the baseline.
Codes:
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-4, R-1, S-3
(D) P-3, Q-1, R-4, S-2

NODIA

(A) Functional Cohesion
(B) Temporal Cohesion
(C) Coincidental Cohesion
(D) Logical Cohesion

183. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A good software design exhibits high cohesion.
Reason (R) : High cohesion means that the elements within a module are strongly related and focused on a single task, which makes the module easier to understand, maintain, and reuse.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

184. The activities of Quality Control (QC) are primarily focused on:
(A) Process definition and improvement.
(B) Product-oriented inspection and testing to find defects.
(C) Project planning and risk management.
(D) Auditing and training.

185. A software sizing technique that involves using historical data from past projects to estimate the size of a new, similar project is known as:
(A) Function Point Analysis
(B) Estimation by Analogy
(C) COCOMO
(D) Parkinson's Law

186. Match the Testing Levels in List I with their Primary Focus in List II:
List I: (Testing Level)
P. Unit Testing
Q. Integration Testing
R. System Testing
S. Acceptance Testing
List II: (Primary Focus)
1. Verifying the entire system against its requirements.
2. Validating the software against user needs and business requirements.
3. Testing individual functions or procedures.
4. Testing the interfaces between modules.
Codes:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-1, S-2
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-3, R-1, S-4

187. The process of tracking and controlling the changes in the software is known as:
(A) Version Control
(B) Software Reuse
(C) Reverse Engineering
(D) Change Control

188. Which of the following is NOT a phase of the generic process model for software engineering?
(A) Communication
(B) Planning
(C) Deployment
(D) Financing

189. A non-functional requirement describes:
(A) A specific feature the system must have.
(B) A function that the system should not perform.
(C) A quality attribute or constraint, such as performance or security.
(D) An error in the system's functionality.

190. What is the weakest form of cohesion?
(A) Functional Cohesion
(B) Sequential Cohesion
(C) Coincidental Cohesion
(D) Temporal Cohesion

191. Which of the following is NOT a McCall's Quality Factor?
(A) Integrity
(B) Flexibility
(C) Testability
(D) Modularity

192. Which of the following best describes software sizing?
(A) Estimating the physical memory the software will occupy.
(B) Quantifying the product size as an input to effort and cost estimation.
(C) Choosing the number of developers for a project.
(D) Measuring the screen size for the user

Downloaded from www.pdf.tube

Direction : Please read the following source carefully.

A software process provides a structured framework for the activities, actions, and tasks required to build high-quality software. The Generic Process Model serves as a foundational template, encompassing a set of framework activities such as communication, planning, modeling, construction, and deployment. Each activity is populated by a collection of task sets, which detail the specific work to be done, the deliverables to be produced, and the quality assurance points needed. Process patterns offer a method for describing and solving common problems within the software process, providing a template for a set of activities and tasks. Prescriptive Process Models, such as the Waterfall or Incremental models, advocate for a structured and sequential or iterative approach, defining a predetermined set of process elements and a predictable workflow. In contrast, specialized models like Component-Based Development (CBD) focus on the assembly of pre-existing software components, while Aspect-Oriented Software Development (AOSD) provides mechanisms for modularizing cross-cutting concerns that span multiple components. Formal Methods employ rigorous mathematical notations to specify and verify system properties, offering the highest degree of software assurance. This contrasts with more empirical models like Web Engineering, which adapts software engineering principles to the unique challenges of developing complex, web-based systems.

Based on the above source, attempt the following five questions.

251. According to the Generic Process Model, which of the following best describes the relationship between framework activities and task sets?
(A) Task sets are high-level phases, while framework activities are the specific actions within them.
(B) Framework activities are overarching phases, and each is composed of specific task sets.
(C) Task sets and framework activities are interchangeable terms for the same process elements.
(D) Framework activities are used exclusively in Prescriptive Models, whereas task sets are used in Agile Models.

252. Which process model is described as focusing on the assembly of pre-existing software elements?
(A) Component-Based Development
(B) Aspect-Oriented Software Development
(C) Formal Methods
(D) Web Engineering

253. What is the primary purpose of a process pattern within a software process?
(A) To enforce a strict, sequential workflow identical to the Waterfall model.
(B) To mathematically verify the correctness of the software before construction begins.
(C) To provide a template for solving a commonly occurring problem during development.
(D) To replace the need for generic framework activities like communication and planning.

254. Based on the description, what is a key differentiating characteristic of Formal Methods compared to other process models?
(A) Their focus on iterative development and rapid feedback from stakeholders.
(B) Their reliance on assembling pre-existing components to build new systems.
(C) Their adaptation of principles specifically for the construction of web-based applications.
(D) Their use of mathematical notations for system specification and verification.

255. How do Prescriptive Process Models differ from the general concept of a software process?
(A) They are less structured and allow for more flexibility than a generic process framework.
(B) They focus exclusively on modularizing cross-cutting concerns across the application.
(C) They define a specific, predetermined workflow and set of process elements.
(D) They are only applicable to Web Engineering and not to traditional software systems.

Direction : Please read the following source carefully.

Agile Process Models represent a significant philosophical shift from traditional, prescriptive methodologies. Instead of rigid, long-term planning, they prioritize adaptability, customer collaboration, and the incremental delivery of

201. Which agile methodology is structured around five core activities: Develop an overall model, Build a feature list, Plan by feature, Design by feature, and Build by feature?
(A) Scrum
(B) Extreme Programming (XP)
(C) Feature Driven Development (FDD)
(D) Crystal

202. The document that formally describes the essential requirements of a software product and serves as a contract is the:
(A) Project Plan
(B) Design Document
(C) Software Requirements Specification (SRS)
(D) Test Plan

203. Which of the following design characteristics is most closely associated with the concepts of cohesion and coupling?
(A) Architectural Patterns
(B) User Interface Design
(C) Functional Independence
(D) Data Design

204. Which of the following statements about software quality is incorrect?
(A) Quality Control focuses on defect identification in the product.
(B) Quality Assurance focuses on defect prevention in the process.
(C) McCall's and ISO 9126 are models for defining software quality factors.
(D) High quality is only achieved by testing at the end of the cycle.

205. If project risk is managed by absorbing the potential loss, it is known as:
(A) Risk avoidance
(B) Risk transfer
(C) Risk acceptance or retention
(D) Risk reduction

206. Consider the following statements regarding software testing:
Statement I : Basis path testing is a black-box testing technique.
Statement II : It uses the cyclomatic complexity of a program to determine the number of independent paths.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

207. Which of the following is a primary motivation for software re-engineering?
(A) To add new features to a well-maintained, modern system.
(B) To improve maintainability and understandability of a legacy system.
(C) To discard an old system and build a new one from scratch.
(D) To test the performance of a system under load.

208. In Component-Based Development, components are often commercial off-the-shelf (COTS) entities. What is a major challenge in this approach?
(A) It is too slow for modern development.
(B) It does not promote software reuse.
(C) Trusting the quality and reliability of third-party components.
(D) It requires extensive knowledge of assembly language programming.

209. Which of the following is NOT an agile process model?
(A) Scrum
(B) Dynamic System Development Model (DSDM)
(C) Web Engineering
(D) Adaptive Software Development (ASD)

210. A user story is a natural language description of a feature written from the perspective of an end user. Which model heavily relies on this for requirements?
(A) Waterfall Model
(B) Formal Methods
(C) Agile Models (like Scrum and XP)
(D) V-Model

211. Consider the following outcomes of software design:
I. Stakeholder communication
II. System analysis
III. Large-scale reuse
IV. Detailed procedural logic
Which of the above are benefits of good architectural design?

Click Here to Buy book on Amazon

(D) Review and approve changes to the product backlog.

172. Identify the incorrect pair regarding software testing.
(A) White-box Testing - Focuses on internal code structure.
(B) Black-box Testing - Focuses on functional specifications.
(C) Unit Testing - Tests interactions between modules.
(D) Regression Testing - Ensures new changes do not break existing features.

173. The RMMM plan is a key artifact of which process?
(A) Requirements Elicitation
(B) Software Design
(C) Risk Management
(D) Software Testing

174. Which of the following is a prescriptive software process model?
(A) Scrum
(B) Extreme Programming (XP)
(C) Crystal
(D) Waterfall Model

175. Which design principle advocates for minimizing the connections between modules?
(A) Low Coupling
(B) High Cohesion
(C) Information Hiding
(D) Modularity

176. A tester is given a specification that a password field must be between 8 and 16 characters. They design test cases with 7, 8, 16, and 17 characters. This approach is an example of:
(A) Basis Path Testing
(B) Boundary Value Analysis
(C) Random Testing
(D) Stress Testing

177. Consider the following statements regarding software engineering processes:
Statement I : Reverse Engineering is the process of building a new system from scratch.
Statement II : Software Re-engineering involves modifying a system to improve its quality attributes like maintainability.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

178. A project timeline chart that shows the start and finish dates of the terminal elements and summary elements of a project is known as a:
(A) PERT Chart
(B) Gantt Chart
(C) Flowchart
(D) Use Case Diagram

179. Which of the following quality factors from the ISO 9126 standard refers to the capability of the software product to be understood, learned, used, and attractive to the user?
(A) Functionality
(B) Reliability
(C) Usability
(D) Efficiency

NODIA

180. Consider the following statements about agile development:
I. It is an iterative and incremental approach.
II. It emphasizes comprehensive upfront documentation.
III. It promotes face-to-face communication.
IV. It is resistant to changes in requirements.
V. It delivers working software frequently.
Which of the above statements are true?
(A) I, III, and V only
(B) II, IV, and V only
(C) I, II, and III only
(D) I, IV, and V only

181. What is the primary role of the 'Product Owner' in Scrum?
(A) To manage the development team's daily tasks.
(B) To be responsible for the technical implementation of the product.
(C) Representing stakeholders and managing the backlog to maximize value.
(D) To facilitate Scrum events and remove impediments.

182. A module that performs a set of tasks that are related by the fact that they must all be done in the same span of time (e.g., system initialization) has:

221. Which software process model is known for its five framework activities: communication, planning, modeling, construction, and deployment?
(A) The Spiral Model
(B) The Generic Process Model
(C) The Scrum Model
(D) The V-Model

222. The process of gathering software requirements from clients, users, and other stakeholders is called:
(A) Requirements Specification
(B) Requirements Elicitation
(C) Requirements Validation
(D) Requirements Management

223. A software module that contains a series of functions for calculating different types of taxes exhibits which type of cohesion?
(A) Functional Cohesion
(B) Sequential Cohesion
(C) Logical Cohesion
(D) Procedural Cohesion

224. Consider the following activities related to software quality:
I. Code inspection
II. Defining coding standards
III. Unit testing
IV. Process auditing
V. Peer reviews
Which of the above statement(s) is/are correct examples of quality control activities?
(A) I, III, and V only
(B) II and IV only
(C) I, II, and III only
(D) I, II, III, IV, and V

225. What is the primary function of a time-line chart like a Gantt chart?
(A) To model the data flow within a system.
(B) To illustrate the schedule, showing tasks, durations, and dependencies.
(C) To calculate the total cost of the project.
(D) To show the risk probability and impact matrix.

226. Which testing methodology is characterized by having some knowledge of the internal workings of the system?
(A) Black-box testing
(B) Grey-box testing
(C) Alpha testing
(D) Beta testing

227. What is the main difference between software re-engineering and reverse engineering?
(A) They are identical processes.
(B) Reverse engineering understands the system, while re-engineering alters it.
(C) Re-engineering understands the system, while reverse engineering alters it.
(D) Reverse engineering is part of testing, while re-engineering is part of design.

228. The XP practice of releasing the system in small, frequent increments to the customer is known as:
(A) Pair Programming
(B) Test-Driven Development (TDD)
(C) Small Releases
(D) Continuous Integration

229. A key document produced during the Requirements Engineering phase is the SRS. What does SRS stand for?
(A) System Requirement Standard
(B) Software Requirements Specification
(C) Software Resource Schedule
(D) System Resource Specification

NODIA

230. A design that is easy to maintain and extend is characterized by:
(A) Low cohesion and high coupling
(B) High cohesion and high coupling
(C) Low cohesion and low coupling
(D) High cohesion and low coupling

231. The measure of how well a software system's performance meets its requirements is a part of which ISO 9126 quality factor?
(A) Functionality
(B) Efficiency
(C) Reliability
(D) Usability

232. Which of the following is an example of a project risk?
(A) The software fails to meet performance requirements.
(B) The project exceeds its budget and schedule.

(A) I and IV only
(B) II and III only
(C) I, II, and III only
(D) I, II, III, and IV

212. A company maintains a library of tested and certified software components (e.g., UI widgets) used across projects. This practice is a direct application of:
(A) Software Re-engineering
(B) Reverse Engineering
(C) Software Reuse
(D) Risk Mitigation

213. Arrange the following testing techniques from the one requiring the least knowledge of the system's internal workings to the one requiring the most:
I. White-box testing
II. Black-box testing
III. Grey-box testing
(A) I, III, II
(B) II, I, III
(C) II, III, I
(D) III, II, I

214. Which of the following is considered a Framework Activity in the generic software process model?
(A) Coding a specific module
(B) Performing a code review
(C) Construction
(D) Creating a test case

215. Consider the following statements regarding software requirements:
I. Functional requirements define what the system does.
II. Non-functional requirements define how well the system does it.
III. The system shall be written in Java is a functional requirement.
IV. A use case is a tool for capturing non-functional requirements.
Which combination of the above statements is correct?
(A) I and II
(B) I and III
(C) II and IV
(D) III and IV

216. The design of classes, their attributes, methods, and the relationships among them is a central part of:
(A) Data Design
(B) Architectural Design
(C) Object-Oriented Design
(D) User Interface Design

217. In the context of the ISO 9126 quality model, the capability of the software product to be modified falls under which characteristic?
(A) Reliability
(B) Usability
(C) Maintainability
(D) Portability

218. In a COCOMO estimation, if a project involves developing complex control systems or real-time software, it would most likely be classified under which mode?
(A) Organic
(B) Semi-detached
(C) Embedded
(D) Simple

219. A tester finds that after a bug fix for the login page, the user profile page now fails to load. To catch such a defect, which type of testing should have been performed?
(A) Unit Testing
(B) Performance Testing
(C) Regression Testing
(D) Stress Testing

220. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Change control is often a bureaucratic and slow process.
Reason (R) : Procedures ensure changes are evaluated for impact and approved systematically, maintaining the integrity of the project baseline.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

Downloaded from www.pdf.tube

(C) Subjecting the system to extreme loads to observe failure.
(D) Testing individual components in isolation.

241. Software reuse is most effectively implemented through which of the following approaches?
(A) Copying and pasting code between projects.
(B) Using standardized, documented, and componentized assets.
(C) Re-writing all code from scratch for each new project.
(D) Only hiring developers who have worked on similar projects.

242. The spiral model of software development combines elements of which two other models?
(A) Waterfall and Scrum
(B) Prototyping and Waterfall
(C) XP and Scrum
(D) FDD and Crystal

243. Which of the following is the best example of a non-functional requirement?
(A) The system must allow users to reset their password.
(B) The system must authenticate users before granting access.
(C) The system must process a transaction in under 3 seconds.
(D) The system must generate a monthly sales report.

244. In software design, if the elements of a module are grouped together because they are executed in a specific order, what type of cohesion is this?
(A) Temporal Cohesion
(B) Logical Cohesion
(C) Sequential Cohesion
(D) Communicational Cohesion

245. The RMMM plan is an output of which activity?
(A) Quality Assurance
(B) Software Testing
(C) Risk Management
(D) Requirements Analysis

246. What is the primary purpose of a time-line chart in project management?
(A) To allocate budget to different tasks.
(B) To visually represent the schedule and track progress against it.
(C) To document the software requirements.
(D) To identify and mitigate project risks.

247. The practice of running tests for a feature before the feature is written is known as:
(A) Regression Testing
(B) Acceptance Testing
(C) Performance Testing
(D) Test-Driven Development (TDD)

248. Arrange the following generic process framework activities in their logical order:
I. Construction
II. Planning
III. Deployment
IV. Communication
V. Modeling
(A) IV, II, V, I, III
(B) I, II, III, IV, V
(C) IV, V, II, I, III
(D) II, IV, V, I, III

NODIA

249. Consider the following software design concepts:
I. Abstraction
II. Cohesion and Coupling
III. Modularity
IV. Information Hiding
V. Requirement Elicitation
Which of the above statements refer to key design concepts in software engineering?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, and V only
(D) I, IV, and V only

250. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Verification and Validation are both necessary for ensuring software quality.
Reason (R) : Verification ensures the software is built correctly according to specifications, while Validation ensures the right software is built to meet user needs.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

(C) The software has a critical security vulnerability.
(D) A competitor releases a superior product.

233. Consider the following statements regarding software testing:
Statement I : Unit testing is a white-box testing technique.
Statement II : Integration testing can use both black-box and white-box techniques.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is false, but Statement II is true.
(D) Statement I is true, but Statement II is false.

234. In version control systems, what is the purpose of branching?
(A) To delete a version of the software permanently.
(B) To create a separate line of development for new features.
(C) To merge two different versions of the software.
(D) To revert the code to a previous state.

235. The Manifesto for Agile Software Development was created as a response to the perceived overhead of which type of process model?
(A) Component-based models
(B) Formal method models
(C) Prescriptive, heavyweight models like Waterfall
(D) Other agile models

236. Match the following Terms in List I with their Definitions in List II:
List I: (Term)
P. Error
Q. Fault
R. Failure
S. Bug
List II: (Definition)
1. An incorrect step, process, or data definition by a person.
2. The inability of a system to perform its required function.
3. A static defect in the software.
4. A common industry term for a fault.
Codes:
(A) P-1, Q-3, R-2, S-4
(B) P-3, Q-1, R-2, S-4
(C) P-1, Q-2, R-3, S-4
(D) P-2, Q-3, R-1, S-4

237. Which of the following is a key characteristic of the Dynamic System Development Model (DSDM)?
(A) It is best suited for projects with stable, unchanging requirements.
(B) It uses timeboxing to fix time and resources, adjusting functionality.
(C) It requires extensive documentation before any development begins.
(D) It forbids any direct user involvement during development.

238. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Information hiding is a fundamental principle of good software design.
Reason (R) : It reduces module coupling by hiding the internal implementation details of a module, making the system more robust to changes.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

239. Consider the following software lifecycle activities:
I. Project Scheduling
II. Risk Management
III. Cost Estimation
IV. Unit Testing
V. Requirements Elicitation
Which of the above statement(s) is/are associated with software project management?
(A) I, II, and V only
(B) I, II, and III only
(C) III, IV, and V only
(D) I, IV, and V only

240. Which of the following best describes Stress Testing?
(A) Verifying that the system meets its functional requirements.
(B) Checking that bug fixes have not introduced new problems.

Click Here to Buy book on Amazon

working software. Extreme Programming (XP) is a prominent agile model characterized by a set of core practices, including pair programming, test-driven development, continuous integration, and frequent releases. It emphasizes technical excellence and a sustainable pace for the development team. Another widely adopted model is Scrum, which organizes work into fixed-length iterations called Sprints. The Scrum framework is defined by specific roles (Product Owner, Scrum Master, Development Team), events (Sprint Planning, Daily Scrum, Sprint Review), and artifacts (Product Backlog, Sprint Backlog). Adaptive Software Development (ASD) is built on a cycle of speculate, collaborate, and learn, acknowledging the uncertainty inherent in complex projects. The Dynamic System Development Model (DSDM) focuses on delivering business value early through timeboxed development and stakeholder involvement. Feature Driven Development (FDD) is a model-driven, short-iteration process that organizes development around specific, client-valued features. Finally, the Crystal family of methodologies suggests that different projects require different policies and practices, tailoring the process to the specific context of the project team and environment.

Based on the above source, attempt the following five questions.

256. Which agile model is explicitly described as organizing work into fixed-length iterations known as Sprints?
(A) Extreme Programming (XP)
(B) Scrum
(C) Feature Driven Development (FDD)
(D) Adaptive Software Development (ASD)

257. What is the core principle of the Crystal family of methodologies as described in the text?
(A) Adherence to a strict set of practices like pair programming and continuous integration.
(B) Organization of development around a comprehensive list of client-valued features.
(C) A focus on the speculative nature of projects through a cycle of learning and collaboration.
(D) Tailoring of development processes to suit the specific context of the project and team.

258. Which of the following is identified as a core practice of Extreme Programming (XP)?
(A) Employing techniques such as pair programming and continuous integration.
(B) Organizing work using a Product Backlog and Sprint Backlog.
(C) Following a continuous cycle of speculate, collaborate, and learn.
(D) Focusing primarily on timeboxed development to deliver business value.

259. Based on the source, what is a common emphasis shared across the described Agile Process Models?
(A) A reliance on formal mathematical methods for verification.
(B) The creation of comprehensive documentation before development begins.
(C) Incremental delivery of software and adaptability to change.
(D) A rigid, sequential progression through predefined development phases.

NODIA

260. The Adaptive Software Development (ASD) model is characterized by a cycle of:
(A) Speculate, Collaborate, Learn.
(B) Plan, Design, Code, Test.
(C) Define, Measure, Analyze, Improve.
(D) Communication, Planning, Modeling, Construction.

Direction : Please read the following source carefully.

The process of defining software requirements is fundamental to the success of any project. Requirements are broadly categorized into functional and non-functional types. Functional requirements specify the behaviors the system must exhibit, defining what the system should do, such as processing a transaction or generating a report. In contrast, non-functional requirements dictate the quality attributes and constraints of the system, describing how the system should perform its functions. Examples include performance, security, reliability, and usability. The initial phase involves eliciting these requirements from stakeholders through various techniques. A key method for capturing functional requirements from a user's perspective is developing use cases, which describe sequences of interactions between an actor (a user or another system) and the

system to achieve a specific goal. Subsequently, requirement analysis and modeling are performed to refine understanding, resolve inconsistencies, and build a conceptual model of the requirements. A requirements review is a formal quality assurance activity where stakeholders and the development team scrutinize the requirements for clarity, completeness, and feasibility. The entire process culminates in the creation of the Software Requirement and Specification (SRS) document, which serves as an official, detailed contract of what the system will do, forming the basis for design, implementation, and verification.

Based on the above source, attempt the following five questions.

261. A requirement stating that the system must process 1000 transactions per second would be classified as which type?
(A) A functional requirement, as it defines a system process.
(B) A use case, as it involves system interaction.
(C) An elicited requirement, as it comes from a stakeholder.
(D) A non-functional requirement, as it specifies a performance constraint.

262. What is the primary purpose of developing use cases during the requirements phase?
(A) To define the quality attributes and performance constraints of the system.
(B) To describe sequences of interactions between an actor and the system.
(C) To serve as the final, formal contract for the software development.
(D) To model the internal data structures and architecture of the system.

263. According to the text, what is the ultimate output of the software requirements process?
(A) A set of unresolved stakeholder needs.
(B) A collection of individual use cases.
(C) The Software Requirement and Specification (SRS) document.
(D) A successful requirements review meeting.

264. What role does a requirements review play in the software development process?
(A) It is a coding activity to implement the specified requirements.
(B) It is a final user acceptance test performed after deployment.
(C) It is an informal discussion for brainstorming potential system features.
(D) It is a quality assurance activity to check requirements for clarity and completeness.

265. Which statement accurately distinguishes between functional and non-functional requirements?
(A) Functional requirements describe system behaviors, while non-functional requirements describe quality attributes.
(B) Functional requirements are elicited from users, while non-functional requirements are defined by developers.
(C) Functional requirements are documented in the SRS, while non-functional requirements are not.
(D) Functional requirements are easy to test, while non-functional requirements are impossible to test.

Direction : Please read the following source carefully.

Effective software design is guided by a set of fundamental principles that promote quality, maintainability, and reusability. Abstraction is a key concept that involves simplifying complex reality by modeling classes, objects, and behaviors relevant to the problem, while hiding irrelevant implementation details. This is closely related to Information Hiding, a principle that suggests modules should be designed so that the information contained within a module (e.g., data structures, algorithms) is inaccessible to other modules that have no need for it. Modularity and Separation of Concerns are intertwined principles that advocate for partitioning a software system into distinct, independent components, or modules, each addressing a specific aspect of the problem. This division helps manage complexity and allows for parallel development. The quality of modularity is often measured by the principles of Functional Independence, which is assessed through two qualitative criteria: Cohesion and Coupling. Cohesion refers to the degree to which the elements within a single module are functionally related and work together to provide a well-defined set of tasks. High cohesion is desirable. Coupling, conversely, measures the degree of interdependence between different modules.

Downloaded from www.pdf.tube

Low coupling, indicating minimal dependency between modules, is a hallmark of good design, as it reduces the ripple effect of changes made in one module on others.

Based on the above source, attempt the following five questions.

266. Which design principle advocates for partitioning a system into distinct components that each address a specific aspect of the problem?
(A) Abstraction
(B) Modularity
(C) Information Hiding
(D) Cohesion

267. According to the text, high cohesion is a desirable attribute because it indicates that:
(A) The module is highly dependent on many other modules within the system.
(B) The elements within a module are functionally related and work towards a common purpose.
(C) The internal implementation details of a module are visible to all other modules.
(D) The module simplifies a complex reality by hiding irrelevant details.

268. What is the relationship between Functional Independence, Cohesion, and Coupling?
(A) High cohesion and low coupling are indicators of high functional independence.
(B) High cohesion and high coupling lead to high functional independence.
(C) Low cohesion and low coupling lead to high functional independence.
(D) Functional independence is a measure of abstraction, unrelated to cohesion and coupling.

269. The principle of Information Hiding is primarily concerned with:
(A) Dividing the system into separate, independent components.
(B) Ensuring that all elements within a module are functionally related.
(C) Restricting access to the internal details of a module.
(D) Simplifying complexity by modeling only relevant objects and behaviors.

270. A software design where a change in one module necessitates extensive changes in several other modules is exhibiting:
(A) High cohesion
(B) Low abstraction
(C) Low coupling
(D) High coupling

NODIA

Direction : Please read the following source carefully.

Software design is a multi-faceted process that translates user requirements into a blueprint for constructing the software. This process operates at different levels of abstraction. At the highest level is Architectural Design, which establishes the overall structure of the system. It defines the major software components, their properties, and the relationships and interactions among them, thereby setting the foundational framework for the entire system. Following this, Data Design focuses on creating the data models and structures that the system will use. This involves transforming the data objects identified during analysis into database schemas and defining the attributes and relationships of the data that the software will manipulate. A critical aspect of modern software is User Interface (UI) Design, which deals with the look and feel of the system and the design of the interaction between the user and the software. It aims to create an intuitive, efficient, and aesthetically pleasing interface. Finally, at the lowest level of detail is Component Level Design. This activity focuses on the internal workings of each individual component or module identified in the architectural design. It involves developing the detailed procedural logic, algorithms, and data structures necessary for each component to fulfill its specified function, effectively creating a detailed specification for the source code.

Based on the above source, attempt the following five questions.

271. Which design activity is responsible for defining the major software components and their interactions?
(A) User Interface Design
(B) Data Design
(C) Architectural Design
(D) Component Level Design

a modification to the software to ensure that the changes have not adversely affected existing functionality. Performance Testing and Stress Testing are non-functional tests that evaluate the system's responsiveness, stability, and resource consumption under various workloads, particularly at or beyond its limits. Parallel to testing, Software Configuration Management (SCM) is a vital discipline that manages evolution. Its core functions, Change Control and Version Control, provide a mechanism for tracking and controlling modifications to the software, which is essential for maintaining integrity, especially when conducting regression tests or engaging in Software Reuse, Re-engineering, or Reverse Engineering.

Based on the above source, attempt the following five questions.

296. A testing approach where test cases are derived based on the program's control flow graph is an example of:
(A) Black-box testing
(B) Regression testing
(C) Stress testing
(D) White-box testing

297. What is the primary purpose of Regression Testing?
(A) To ensure that modifications have not introduced new defects in existing features.
(B) To evaluate the system's responsiveness and stability under heavy load.
(C) To test the software's functionality without knowledge of its internal code.
(D) To execute every independent path through the program's logic.

298. Which statement accurately distinguishes Black-box testing from White-box testing?
(A) Black-box testing requires knowledge of the code, while White-box testing does not.
(B) Black-box testing focuses on functionality, while White-box testing focuses on internal structure.
(C) Black-box testing is a form of performance testing, while White-box testing is a form of stress testing.
(D) Black-box testing uses Basis Path Testing, while White-box testing uses requirement specifications.

299. According to the source, what are the core functions of Software Configuration Management (SCM)?
(A) Performance Testing and Stress Testing
(B) Change Control and Version Control
(C) White-box and Black-box Testing
(D) Software Reuse and Reverse Engineering

300. Testing a system to determine its stability when subjected to workloads beyond its specified limits is known as:
(A) Stress Testing
(B) Basis Path Testing
(C) Regression Testing
(D) Black-box Testing

www.nodia.in

Click Here to Buy book on Amazon

(C) Usability
(D) Maintainability

279. The activity of conducting a code review to find errors would be best classified as:
(A) Quality Assurance (QA)
(B) Product Transition
(C) Quality Control (QC)
(D) An ISO 9126 characteristic

280. How does the source define Software Reliability?
(A) The ease with which software can be modified to correct faults or adapt to new requirements.
(B) The degree to which a system can be transferred from one hardware or software environment to another.
(C) The effort needed to use the software by a specified set of users.
(D) The probability of failure-free operation for a specified period in a specified environment.

Direction : Please read the following source carefully.

Proactive Risk Management is an essential discipline in software project management, aimed at identifying potential problems before they occur so that risk-handling activities can be planned and invoked. The process begins with risk identification, where potential project, product, and business risks are systematically catalogued. Once identified, risks are subjected to analysis to estimate their probability of occurrence and potential impact or loss. This analysis helps in prioritizing the risks for further attention. The core of risk management is the development and execution of a Risk Mitigation, Monitoring, and Management (RMMM) plan. Risk mitigation involves creating strategies to reduce the likelihood or impact of a risk. Common strategies include avoiding the risk, transferring it to another party, or reducing its effects. Risk monitoring is the ongoing process of tracking the identified risks, assessing their status, and watching for triggers that would indicate a risk is becoming a reality. This allows the project manager to deploy pre-defined contingency plans. Risk management is the overarching set of activities that includes planning, executing, and controlling the mitigation and monitoring efforts throughout the project lifecycle. A key quality attribute often treated as a risk is Software Reliability. A project plan that fails to address the risk of low reliability may lead to a product that is functionally correct but unusable in a real-world operational context.

Based on the above source, attempt the following five questions.

281. What are the three primary activities of the RMMM plan as described in the text?
(A) Risk Mitigation, Risk Monitoring, and Risk Management.
(B) Risk Identification, Risk Analysis, and Risk Prioritization.
(C) Risk Avoidance, Risk Transfer, and Risk Reduction.
(D) Project Risk, Product Risk, and Business Risk.

282. The process of estimating the probability of a risk's occurrence and its potential impact is known as:
(A) Risk identification
(B) Risk analysis
(C) Risk monitoring
(D) Risk mitigation

283. Which of the following best describes the purpose of risk monitoring?
(A) To create a comprehensive list of all potential problems at the start of a project.
(B) To track the status of identified risks and watch for triggers to invoke contingency plans.
(C) To develop strategies that reduce the likelihood and impact of identified risks.
(D) To transfer the responsibility for a specific risk to a third party.

284. A strategy that involves developing a contingency plan to be deployed if a risk becomes a reality is an example of:
(A) Risk mitigation
(B) Risk identification
(C) Risk analysis
(D) Risk cataloguing

285. According to the source, why is it important to manage Software Reliability as a risk?
(A) Because low reliability is the only risk that can affect a software project.
(B) Because reliability can only be managed by

Downloaded from www.pdf.tube

assets rather than building every feature from scratch. It promotes efficiency and reliability, as components are typically tested and maintained independently before being integrated into larger, more complex software applications.

9. (C) I, II, and III only
The generic software process model incorporates essential framework activities such as planning, modeling, construction, and deployment. These phases represent the fundamental stages of engineering, from initial strategy to final release. Non-engineering tasks like marketing or ongoing user training are generally excluded from this specific process framework. Success depends on the disciplined execution of these technical actions to deliver quality products.

10. (A) The Waterfall Model
The waterfall model is a classic prescriptive process that follows a strict linear sequence of phases. In this approach, each stage must be completed before the next begins, with the output of one phase serving as the direct input for the following one. This sequential structure provides high levels of discipline but offers limited flexibility when requirements change mid-project.

11. (B) I, III, and IV
Agile process models prioritize iterative delivery and high adaptability to changing requirements throughout the software lifecycle. Unlike traditional prescriptive models, agile methodologies value customer collaboration over rigid contracts and favor working software over extensive documentation. This flexible approach allows teams to respond quickly to feedback, ensuring the final product aligns closely with the evolving needs of the end users.

12. (B) To use mathematical specifications for designing and verifying a system.
Formal methods involve the use of mathematical specifications to design and verify software systems rigorously. This approach allows developers to prove the correctness of a system through logical reasoning and mathematical proofs. By eliminating ambiguity and detecting errors early, formal methods are particularly valuable in developing safety-critical systems where reliability and precision are paramount for project success and safety.

13. (C) Scrum
Scrum is an agile framework that organizes work into fixed-length iterations called sprints. These cycles typically last two to four weeks and aim to produce a potentially shippable product increment. This iterative structure allows teams to incorporate feedback regularly and maintain a steady pace of delivery. It fosters transparency, inspection, and adaptation, which are the core pillars of Scrum.

14. (D) The 80/20 rule, where 80% of the application can be delivered in 20% of the time.
The Dynamic System Development Model is an agile methodology that operates on the principle that eighty percent of an application can be delivered in twenty percent of the time. It focuses on rapid delivery and active user involvement to ensure the system meets business needs. DSDM provides a disciplined framework for agile projects, balancing speed with quality control.

15. (A) To facilitate the Scrum process and remove impediments for the development team.
In the Scrum framework, the Scrum Master serves as a facilitator who ensures the team adheres to Scrum practices and values. Their primary responsibility is to remove any impediments that hinder the development team's progress during a sprint. Rather than managing the project budget or technical design, they focus on optimizing team performance and fostering a collaborative, self-organizing environment.

16. (B) Comprehensive Documentation
Extreme Programming is built upon core values including communication, simplicity, feedback, and courage. These values guide technical practices like pair programming and continuous integration to improve software quality. Unlike traditional models, XP favors working software and responsiveness to change over comprehensive documentation. This approach ensures that the development process remains streamlined and focused on delivering real value to stakeholders.

17. (A) I, II, III, IV, V
The generic waterfall model follows a sequential progression beginning with system engineering

NODIA

given development phase satisfy the conditions imposed at the start of that phase; it is often summarized as are we building the product right?. Validation is the process of evaluating software during or at the end of the development process to determine whether it satisfies specified business requirements; it is summarized as are we building the right product?. These processes aim to uncover defects, which originate from a human Error that results in a Fault (or Bug) in the code, which in turn may cause a Failure, an observable deviation from the expected behavior of the system. Testing is conducted at different levels. Unit Testing focuses on verifying the functionality of the smallest individual components or modules of the software in isolation. Following this, Integration Testing is performed to expose defects in the interfaces and interactions between integrated components. Towards the end of the development cycle, forms of user acceptance testing such as Alpha Testing (conducted by internal teams at the developer's site) and Beta Testing (conducted by real users in their own environment) are used to validate the software against user needs before release.

Based on the above source, attempt the following five questions.

291. Which of the following questions best represents the concept of Verification?
(A) Does the software meet the user's business needs?
(B) Are we building the right product?
(C) Does the software conform to its design specification?
(D) Is the product ready for release to the general public?

292. According to the text, what is the correct causal relationship between Error, Fault, and Failure?
(A) An Error leads to a Fault, which may cause a Failure.
(B) A Fault leads to an Error, which causes a Failure.
(C) A Failure leads to a Fault, which is caused by an Error.
(D) A Fault, an Error, and a Failure are all synonymous terms.

293. Testing individual software modules in isolation to check their functionality is known as:
(A) Integration Testing
(B) Beta Testing
(C) Validation
(D) Unit Testing

294. What is the key difference between Alpha Testing and Beta Testing?
(A) Alpha testing is performed by real users, while beta testing is performed by internal teams.
(B) Alpha testing is done at the developer's site, while beta testing is done at the end-user's site.
(C) Alpha testing focuses on individual units, while beta testing focuses on integrated components.
(D) Alpha testing is verification, while beta testing is validation.

295. The primary goal of Integration Testing is to:
(A) Verify that the software meets all business requirements.
(B) Test individual components in a completely isolated environment.
(C) Expose defects in the interfaces and interactions between components.
(D) Validate the software with a select group of external end-users.

Direction : Please read the following source carefully.

Software testing encompasses a variety of techniques designed to systematically uncover different types of defects. These techniques are often categorized as White-box and Black-box testing. White-box testing derives test cases from the internal structure and logic of the software. Testers use their knowledge of the code to ensure that all internal paths and conditions are exercised. Specific white-box techniques include Basis Path Testing, which aims to execute every independent path through a program's control flow graph, and Control Structure Testing, which includes condition and loop testing. In contrast, Black-box testing focuses solely on the functional requirements of the software, deriving test cases from the specification without any knowledge of the internal implementation. It treats the software as an opaque box, testing its inputs and verifying the corresponding outputs. Beyond these fundamental approaches, specialized testing is crucial. Regression Testing is performed after

Click Here to Buy book on Amazon

transferring the risk to another party.

(C) Because a functionally correct product may be unusable if its reliability is low.

(D) Because reliability risks are identified only after the product has been deployed.

Direction : Please read the following source carefully.

The estimation and scheduling of software projects are critical management activities that determine project feasibility, resource allocation, and delivery timelines. A foundational step is Software Sizing, which quantifies the size of the software to be built. Two common metrics for this are Lines of Code (LOC) and Function Points (FP). LOC is a direct measure of the code volume, while FP is an indirect measure based on the functionality delivered to the user, calculated from components like inputs, outputs, inquiries, files, and interfaces. Once size is estimated, this value is used to estimate the required cost and effort. Various Estimation Models exist for this purpose, with the Constructive Cost Model (COCOMO) being one of the most well-known. COCOMO is an algorithmic model that uses a basic regression formula with parameters derived from historical project data to predict effort and duration based on the estimated size of the software in KLOC (Kilo Lines of Code). Following estimation, Project Scheduling is performed to distribute the estimated effort across the project timeline. This involves defining tasks, identifying dependencies, and assigning personnel. Time-line Charts, such as Gantt charts, are widely used tools for visualizing the project schedule, showing task start and end dates, durations, and dependencies, which aids in tracking progress and managing staffing.

Based on the above source, attempt the following five questions.

286. What is the fundamental difference between Lines of Code (LOC) and Function Point (FP) based estimations?

(A) LOC is used in the COCOMO model, while FP is used exclusively for scheduling.

(B) LOC measures inputs and outputs, while FP measures the number of lines in the final program.

(C) LOC is an algorithmic model, while FP is a visualization tool like a Gantt chart.

(D) LOC is a direct measure of code volume, while FP is an indirect measure based on functionality.

287. The Constructive Cost Model (COCOMO) is primarily used for what purpose?

(A) To visualize the project timeline and task dependencies.

(B) To quantify the functionality delivered to the user.

(C) To predict project effort and duration based on software size.

(D) To manage the allocation of staff to different project tasks.

288. According to the text, what is the primary input for the COCOMO model?

(A) The number of function points.

(B) The estimated size of the software in Kilo Lines of Code.

(C) The project's budget in dollars.

(D) A detailed time-line chart.

289. What is the main function of a Time-line Chart in project management?

(A) To calculate the number of function points for a system.

(B) To estimate the total cost and effort required for the project.

(C) To provide a regression formula for predicting project duration.

(D) To visualize the project schedule, tasks, and their durations.

290. Which sequence of activities is logically correct according to the source?

(A) Software Sizing -> Effort Estimation -> Project Scheduling.

(B) Project Scheduling -> Software Sizing -> Effort Estimation.

(C) Effort Estimation -> Project Scheduling -> Software Sizing.

(D) Project Scheduling -> Effort Estimation -> Software Sizing.

NODIA

Direction : Please read the following source carefully.

Software testing is a critical activity for ensuring quality, involving both Verification and Validation. Verification is the process of evaluating software to determine whether the products of a

ANSWERS

1. (B) A collection of activities and tasks required for creating a quality product.
Software processes encompass the structured set of activities, actions, and tasks required to create high-quality software products. This framework provides a roadmap for development teams by defining the necessary steps for project management, technical work, and product delivery. Following such a process ensures consistency, predictability, and efficiency throughout the entire development lifecycle for teams.

2. (A) Both Statement I and Statement II are true.
Software engineering actions are defined by specific task sets that outline the actual work required to achieve objectives. Simultaneously, framework activities represent high-level stages like communication or deployment that guide the overall process. Both concepts are fundamental to establishing a disciplined approach, ensuring that every phase of development is clearly defined, actionable, and measurable for project success.

3. (C) Marketing
Generic framework activities in software engineering typically include communication, planning, modeling, construction, and deployment. These phases represent the technical and managerial workflow required to build software. Marketing is a business function external to the actual engineering process. It focuses on product promotion and sales rather than the technical activities involved in the systematic development of software applications.

4. (D) A template describing a recurring problem and a proven solution.
Process patterns provide a template for describing a recurring problem encountered during the software process and suggest a proven solution. These patterns help organizations learn from past experiences by formalizing successful approaches to common challenges. They enhance the consistency of software development efforts by providing a reusable framework for addressing issues related to project management or technical execution.

5. (C) They are designed to be highly adaptive to rapid changes in requirements.
Prescriptive process models focus on structure and order by defining a distinct set of elements and a predictable workflow. While models like the waterfall provide stability, they are often criticized for their inability to handle rapid changes in requirements. These models prioritize rigid planning and sequential execution over the high adaptability found in modern agile development methodologies.

6. (A) P-4, Q-3, R-2, S-1
Agile process models offer distinct methodologies for development teams. Extreme Programming emphasizes technical excellence and human values, while Scrum utilizes time-boxed sprints for iterative delivery. Crystal provides a family of methodologies based on project size, and Feature Driven Development organizes work around specific functional features. Each model addresses unique project constraints while maintaining a focus on incremental delivery.

7. (D) Identifying, specifying, and representing cross-cutting concerns.
Aspect-Oriented Software Development focuses on managing cross-cutting concerns that affect multiple parts of a system, such as security or logging. These concerns are often difficult to modularize using traditional object-oriented techniques. By identifying and representing these aspects separately, developers can achieve better modularity and cleaner code. This approach simplifies maintenance and enhances the overall structure of complex software systems.

8. (A) Both A and R are true and R is the correct explanation of A.
Component-based development centers on the construction of software systems using pre-existing, modular components designed for reuse. This model significantly reduces development time and costs by leveraging verified

Downloaded from www.pdf.tube

272. The process of transforming data objects into database schemas is a primary focus of which design type?
(A) Component Level Design
(B) Data Design
(C) Architectural Design
(D) User Interface Design

273. Which of the following best describes Component Level Design?
(A) Establishing the high-level structure and framework of the entire system.
(B) Defining the aesthetic and interaction model for the end-user.
(C) Translating analysis objects into a comprehensive database schema.
(D) Creating the detailed internal logic and algorithms for individual modules.

274. Based on the levels of abstraction described, what is the logical sequence of design activities?
(A) Architectural Design -> Data Design / User Interface Design -> Component Level Design
(B) Component Level Design -> Architectural Design -> User Interface Design
(C) User Interface Design -> Data Design -> Architectural Design
(D) Data Design -> Component Level Design -> Architectural Design

275. The design of the interaction between the user and the software is the main concern of:
(A) Architectural Design
(B) Component Level Design
(C) User Interface Design
(D) Data Design

Direction : Please read the following source carefully.

Defining and measuring software quality is a complex endeavor addressed by several structured models. McCall's Quality Factors model, one of the early frameworks, categorizes software quality attributes into three major perspectives: product operation (correctness, reliability, efficiency), product revision (maintainability, flexibility, testability), and product transition (portability, reusability, interoperability). A more contemporary framework is the ISO 9126 standard, which defines six primary quality characteristics: Functionality, Reliability, Usability, Efficiency, Maintainability, and Portability. Each of these characteristics is further broken down into sub-characteristics to allow for more detailed evaluation. In practice, achieving quality involves two distinct but related activities: Quality Control (QC) and Quality Assurance (QA). Quality Control is a product-oriented set of activities that focuses on defect detection and removal. It includes activities like code reviews and testing to find faults in the finished product. In contrast, Quality Assurance is a process-oriented set of activities designed to prevent defects from occurring in the first place. QA involves defining development standards, conducting process audits, and ensuring that the team follows the established procedures to build quality into the product from the start. Software Reliability, defined as the probability of failure-free operation for a specified time in a specified environment, is a key quantitative metric within both models.

Based on the above source, attempt the following five questions.

276. According to McCall's Quality Factors, the attribute of 'maintainability' falls under which major perspective?
(A) Product revision
(B) Product operation
(C) Product transition
(D) Product deployment

277. What is the primary difference between Quality Assurance (QA) and Quality Control (QC)?
(A) QA is product-oriented, focusing on defect detection; QC is process-oriented, focusing on defect prevention.
(B) QA is performed by the testing team, while QC is performed by the development team.
(C) QA is concerned with McCall's factors, while QC is concerned with ISO 9126 factors.
(D) QA is process-oriented, focusing on defect prevention; QC is product-oriented, focusing on defect detection.

278. Which of the following is NOT one of the six primary quality characteristics defined by the ISO 9126 standard?
(A) Reliability
(B) Cost-effectiveness

NODIA

and modeling, followed by requirements analysis and design. Once the design is finalized, the process moves into code generation and finally testing. This linear approach ensures that each phase has a clearly defined set of goals and deliverables before the project moves to the next developmental stage.

18. (A) Systematic and disciplined approaches to developing high-quality Web systems.
Web engineering applies systematic and disciplined approaches to the development, deployment, and maintenance of high-quality web-based systems. It addresses unique challenges such as concurrency, network-heavy environments, and security risks inherent to the internet. By using specialized methodologies and tools, web engineering ensures that applications are robust, scalable, and capable of meeting the diverse needs of a global user base.

19. (C) Adaptive Software Development (ASD) - An agile model focused on collaboration and learning cycles.
Adaptive Software Development is an agile process model that centers on collaboration and learning cycles. It replaces traditional planning with speculating, collaborating, and learning to manage complex systems. This approach acknowledges the inherent uncertainty in software projects and encourages teams to adapt their strategies based on continuous feedback. It promotes a culture of experimentation and improvement during development.

20. (D) Functional Requirement
Functional requirements specify the services and behaviors the software must provide to its users. In a hospital management system, allowing a doctor to access patient records is a core function that defines what the system does. These requirements are essential for satisfying user needs and are typically documented in detail to guide the subsequent design and development phases.

21. (A) Both A and R are true and R is the correct explanation of A.
Non-functional requirements describe the quality attributes and constraints of a system, such as performance, security, and reliability. They focus on how the system performs its tasks rather than the specific functions it provides. These requirements are critical for user satisfaction and operational success, as they define the technical environment and performance standards necessary for the software to function.

22. (B) Developing Use Cases
Use cases are a requirements elicitation technique that describes the interactions between an external actor and the system to achieve a specific goal. By focusing on user scenarios, they help developers understand the system's functional requirements from the perspective of the end user. This method is effective for identifying boundary conditions and ensuring that the system fulfills user expectations.

23. (A) I and IV
The Software Requirements Specification serves as a formal contract between the client and the developer, outlining all functional and non-functional requirements. A high-quality SRS must be unambiguous, consistent, and complete to prevent misunderstandings during development. It acts as the foundation for the entire project, guiding design, coding, and testing activities while ensuring the final product meets expectations.

24. (A) Non-Functional Requirement
Requirements that define performance metrics, such as handling a specific number of concurrent users or maintaining response times, are classified as non-functional requirements. They impose constraints on the system's operation rather than describing specific features. These attributes are vital for ensuring the system remains usable and efficient under varying workloads, directly impacting the overall quality and user experience.

25. (C) To assess the quality and completeness of the requirements specification.
A requirements review is conducted to assess the quality, clarity, and completeness of the requirements specification. This process involves stakeholders and developers examining the document to identify errors, inconsistencies, or omissions before the design phase begins. By catching issues early, the team can avoid costly rework and ensure that the project is built on a solid, agreed-upon foundation.

Click Here to Buy book on Amazon

product. This process is essential for coordinating team efforts, preventing conflicting changes, and maintaining a clear history of the system's evolution over time.

81. (B) Version Control
Version control is a primary activity of Software Configuration Management that involves tracking and managing multiple versions of configuration items like source code and documentation. It allows teams to collaborate efficiently, revert to previous states if errors occur, and maintain a clear audit trail of changes. Using version control systems is a standard practice for ensuring project stability and organization.

82. (B) Software Re-engineering
Software re-engineering involves creating a new version of a program from an existing one by adding new functionalities or improving quality attributes like maintainability. This process typically includes analyzing the existing system, restructuring the code, and then implementing changes. It is used to modernize legacy systems, making them more efficient and easier to support in evolving technical environments.

83. (A) Reverse Engineering
Reverse engineering is the process of analyzing a subject system to identify its components and their relationships, often creating representations at a higher level of abstraction. It is frequently used to understand legacy code that lacks proper documentation. By extracting design information from the source code, developers can gain insights necessary for maintenance, re-engineering, or reusing existing software assets.

84. (B) Building software systems using existing assets rather than from scratch.
Software reuse involves building new systems by leveraging existing assets, such as code components, designs, or documentation, rather than starting from scratch. This approach increases productivity, reduces development costs, and improves software quality by using previously tested and verified modules. Effective reuse requires a disciplined strategy for identifying, cataloging, and integrating reusable components into the software engineering process.

85. (C) I, II, III, and V only
Core activities of Software Configuration Management include identifying configuration items, establishing baselines, managing changes to those baselines, and reporting status. These tasks ensure that every modification to the system is documented, reviewed, and controlled. While unit testing is a quality control activity, SCM focuses on the structural integrity and historical tracking of the software's components and their various versions.

86. (B) To approve or reject proposed changes to a software baseline.
The Change Control Board is a formal group responsible for reviewing, evaluating, and approving or rejecting proposed changes to a software project's baseline. Their role is to ensure that modifications are necessary, well-understood, and do not negatively impact other parts of the system or the project schedule. The CCB provides a centralized authority for managing change systematically and transparently.

87. (C) Waterfall Model
The waterfall model is most appropriate for projects where requirements are well-understood, stable, and unlikely to change significantly. Its structured, linear approach provides a clear roadmap for development, making it easy to manage and track progress. In environments where a formal, step-by-step process is preferred, the waterfall model offers a predictable framework for delivering a finalized software product.

88. (C) It requires a detailed, upfront design before any coding begins.
Agile development emphasizes working software and adaptability over comprehensive documentation and rigid upfront planning. Unlike traditional models, agile does not require a detailed design to be finalized before coding begins. Instead, it promotes sustainable development and welcomes changing requirements throughout the project. This allows teams to deliver value incrementally and respond effectively to user feedback as the product evolves.

89. (B) Scrum
The team structure consisting of a Product Owner, a Scrum Master, and a Development

NODIA

combination results in software that is robust, flexible, and maintainable.

35. (C) Functional Cohesion
Functional cohesion occurs when all elements within a module contribute to the performance of a single, well-defined task. This is considered the strongest and most desirable form of cohesion in software design. Modules with high functional cohesion are easier to understand, test, and reuse because their purpose is clear and they do not have side effects on unrelated functionalities.

36. (A) Both A and R are true and R is the correct explanation of A.
Low coupling is a desirable characteristic because it minimizes the interdependence between software modules. This reduces the ripple effect, where a change in one part of the system necessitates changes or causes errors in other parts. By isolating modules, low coupling makes the software more resilient to change, simplifies unit testing, and enhances the overall maintainability and stability of the system.

37. (C) To describe the system structure, components, and relationships.
Architectural design focuses on describing the high-level structure of a system, its components, and the relationships between them. It provides a blueprint for the entire project, ensuring that the software meets its quality attributes and requirements. By establishing a solid architecture, developers can manage complexity and facilitate communication among stakeholders while providing a framework for more detailed design activities.

38. (B) Modules should be designed so that internal information is inaccessible to those without need.
Information hiding suggests that modules should be designed so that their internal data and implementation details are inaccessible to other parts of the system. This principle reduces dependencies between modules by exposing only the necessary interface for interaction. It enhances security and maintainability, as changes to a module's internal logic do not affect other components as long as the interface remains consistent.

39. (D) It transforms structural elements into procedural descriptions.
Component-level design involves transforming structural elements into detailed procedural descriptions of the software components. This phase focuses on the internal logic, algorithms, and data structures required to implement the functionality defined in higher-level design documents. By specifying these details, developers create a clear path for coding, ensuring that each component performs its intended tasks accurately and efficiently.

40. (A) Modularity - The process of hiding the internal details of a module.
Modularity refers to the process of decomposing a complex system into smaller, manageable, and self-contained units. It is not about hiding internal details, which is the definition of information hiding. High cohesion and low coupling are measures used to evaluate the quality of modularity. Abstraction and information hiding are related but distinct concepts that complement modularity in creating effective software designs.

NODIA

41. (D) Modularity
The decomposition of a banking application into separate modules for customer management, loan processing, and transactions is an application of the modularity principle. By breaking the system into manageable and specialized units, the design becomes easier to develop, test, and maintain. This approach also allows different teams to work on various parts of the system simultaneously, improving overall project productivity.

42. (B) To make user interaction as simple and efficient as possible.
A primary goal of user interface design is to make user interaction as simple, efficient, and intuitive as possible. Good UI design focuses on the user's needs and tasks, ensuring that the interface is easy to navigate and provides clear feedback. By reducing complexity and cognitive load, developers can enhance the overall user experience and improve the effectiveness of the software.

43. (C) Usability
In McCall's quality model, usability refers to the effort required for a user to learn, operate, and interpret the output of a software program.

26. (B) P-3, Q-4, R-1, S-2
Various techniques are used to gather requirements effectively from stakeholders. Interviews provide direct interaction for deep insights, while use cases capture system behavior from a user's perspective. Prototypes offer a working model for immediate feedback, and questionnaires allow for data collection from a large audience. Using a combination of these methods ensures a comprehensive understanding of the system's requirements.

27. (D) Ambiguous
A high-quality Software Requirements Specification should be verifiable, modifiable, and traceable to ensure the project can be managed effectively. Ambiguity is a negative characteristic that leads to misunderstandings and errors during development. Avoiding vague language ensures that both the developers and the client have a shared, precise understanding of what the software is intended to accomplish throughout its lifecycle.

28. (B) Requirements Engineering
Requirements engineering is the systematic process of discovering, analyzing, documenting, and validating the needs of stakeholders for a software system. This phase is critical because it establishes the baseline for all subsequent development activities. Proper execution ensures that the software correctly addresses the intended problem and provides a clear roadmap for the design, construction, and testing of the product.

29. (A) Both Statement I and Statement II are true.
Functional requirements describe the specific services or functions the software must provide to meet user needs. Conversely, non-functional requirements set constraints on how those services are delivered, focusing on quality attributes like performance and security. Together, these two types of requirements provide a complete picture of the system's expected behavior and operational standards, guiding the engineering process.

30. (D) Hiding implementation details of a module to show functionality.
Abstraction is a design principle that involves hiding the complex implementation details of a module to focus on its high-level functionality. This allows developers to manage complexity by ignoring irrelevant details and concentrating on essential characteristics. By providing a simplified view of components, abstraction facilitates easier understanding, communication, and maintenance within large and complex software systems during the development process.

31. (A) A reusable solution to a common problem within a specific design context.
A design pattern is a reusable solution to a common problem that occurs within a specific context of software design. These patterns provide a standard terminology and a proven approach for developers to address recurring architectural or structural issues. By leveraging design patterns, teams can improve the robustness and maintainability of their software while reducing the time spent on design.

32. (B) Separation of Concerns
Separation of concerns is a fundamental design principle suggesting that a system should be divided into distinct parts with minimal overlap in functionality. Each part, or concern, addresses a specific aspect of the software's overall responsibility. This approach enhances modularity, simplifies debugging, and makes the system easier to maintain, as changes in one area have limited impact on others.

33. (D) I, II, III, and IV
Effective software design relies on several core principles, including modularity, high cohesion, low coupling, and information hiding. These concepts work together to create systems that are easy to understand, test, and maintain. Modularity organizes code into manageable units, while cohesion and coupling ensure those units are focused and independent. Information hiding protects the internal state, further enhancing system stability.

34. (A) Good design aims for high cohesion and low coupling.
In high-quality software design, the goal is to achieve high cohesion and low coupling. High cohesion ensures that the elements within a module are closely related and focused on a single task. Low coupling minimizes the dependencies between different modules, making the system more modular and easier to modify. This

Downloaded from www.pdf.tube

diverse computing landscape, where applications must run on various platforms and devices. High portability ensures that software can be adapted easily, reducing the costs associated with porting code to new systems.

53. (C) Risk Avoidance
Choosing not to include a complex, high-risk feature in an initial release is a classic example of risk avoidance. By eliminating the source of the risk entirely, the team prevents the potential negative impact on the project's schedule or quality. This strategy is often used when the potential consequences of a risk far outweigh the benefits of the risky feature.

54. (B) Software Sizing
Lines of Code is a common metric used for software sizing, providing a quantitative measure of the program's magnitude. It serves as a primary input for many effort and cost estimation models. While easy to count, LOC has limitations, such as being dependent on the programming language and not directly reflecting the complexity or functionality delivered to the end user.

55. (A) I, II, III, IV
Risk management follows a logical sequence beginning with risk identification, where potential threats are listed. Next, risk analysis assesses the probability and impact of these risks. This is followed by risk mitigation planning to develop strategies for addressing them. Finally, risk monitoring ensures that risks are tracked throughout the project lifecycle to respond effectively if they manifest during development.

56. (C) The functionality of the software from the user's point of view.
Function Point analysis is a software sizing method based on the functionality delivered to the user rather than technical implementation details. It measures external inputs, outputs, inquiries, and internal/interface files to estimate the system's size. Because it focuses on the user's perspective, FP analysis is language-independent and can be applied earlier in the development lifecycle than code-based metrics.

57. (C) Estimating project cost and effort
The Constructive Cost Model is an algorithmic approach used for estimating the effort, cost, and schedule of a software project. It uses formulas based on the size of the software and various cost drivers related to the product, hardware, and personnel. COCOMO provides different levels of detail, allowing project managers to refine their estimates as more information becomes available.

NODIA

58. (B) I, II, III, and V only
Valid inputs for calculating Function Points include external inputs, external outputs, external inquiries, internal logical files, and external interface files. These elements represent the functional components of the software from a user's viewpoint. Technical resources, such as the number of developers on a project team, are not used in the size calculation itself, although they impact the final effort estimate.

59. (B) It is language-dependent and can be misleading for measurement.
A major disadvantage of using Lines of Code as a sizing metric is its dependence on the programming language used. This makes it difficult to compare projects written in different languages or to accurately estimate size before coding begins. Additionally, LOC can be misleading because a larger amount of code does not necessarily mean more functionality or higher software quality.

60. (C) Time-line Chart (Gantt Chart)
A Gantt chart is a horizontal timeline chart used in project management to visualize the project schedule. It displays individual tasks, their durations, and their dependencies, providing a clear overview of the project's progress. By showing the start and end dates of all activities, Gantt charts help project managers allocate resources effectively and identify potential delays in the project timeline.

61. (D) Fragmented
The COCOMO model defines three development modes based on project complexity: organic, semi-detached, and embedded. Organic mode is for small, straightforward projects, while embedded mode addresses complex systems with rigid requirements. Semi-detached mode falls in between these two. Fragmented is not a recognized mode in the COCOMO framework, which focuses on the level of technical and organizational difficulty of projects.

62. (A) Both A and R are true and R is the correct explanation of A.
Function Point estimation is often considered more reliable than LOC-based methods because it is independent of the programming language and can be calculated based on requirements early in the lifecycle. This allows for more accurate planning before any code is written. By focusing on functionality rather than syntax, FP provides a consistent measure of software size across different technical environments.

63. (B) Allocate effort and assign tasks to specific team members.
Project scheduling aims to allocate effort and assign specific tasks to team members over the project timeline. This ensures that resources are used efficiently and that the project remains on track to meet its deadlines. A well-defined schedule allows managers to track progress, identify bottlenecks, and communicate expectations clearly, which is vital for the successful delivery of software.

64. (C) COCOMO
COCOMO is a well-known example of an algorithmic cost estimation model, using mathematical formulas to predict project metrics. These formulas are derived from historical data and take into account factors like software size and complexity. Unlike expert judgment or analogy, algorithmic models provide a structured and repeatable way to estimate the resources needed for development based on objective inputs.

65. (A) Both Statement I and Statement II are true.
Effort estimation predicts the amount of labor, usually in person-months, required to complete a project. Cost estimation then calculates the total financial expenditure, which is often derived from the effort estimate plus other overheads. Both are essential for project planning and budgeting, ensuring that the necessary resources and funding are available to support the software development effort.

66. (A) Verification is Are we building the product right?, while Validation is Are we building the right product?.
Verification focuses on whether the software is being built correctly according to the specified requirements, while validation ensures that the final product meets the actual needs of the users. Verification is often described as Are we building the product right?, whereas validation asks Are we building the right product?. Both are essential components of a comprehensive software testing strategy.

67. (C) Error
An error is defined as an incorrect step, process, or data definition in a computer program that leads to unintended behavior. It typically results from a human mistake made during the requirements, design, or coding phases. If left uncorrected, an error can manifest as a fault in the system, which may eventually cause a failure during software execution.

68. (B) II, III, and IV only
In software testing, an error is a human mistake, such as a coding oversight. This error results in a fault, which is a static defect within the code. When this fault is executed, it leads to a failure, which is the observable incorrect behavior of the software. Understanding these distinctions is crucial for effective debugging and quality assurance processes.

69. (B) The tester has full knowledge of the internal structure and logic of the software.
White-box testing is a technique where the tester has full knowledge of the internal structure, logic, and code of the software. This allows them to design test cases that exercise specific paths, branches, and conditions within the program. By focusing on the internal workings, white-box testing helps identify logical errors and ensure that all parts of the code are functioning correctly.

70. (B) White-box testing
White-box testing is also known as glass-box or structural testing because it involves looking inside the box of the software to examine its internal logic. Testers use their knowledge of the source code to verify that all statements and branches are executed as intended. This level of transparency is essential for thorough testing of the system's underlying structure and algorithms.

71. (B) P-1, Q-3, R-2, S-4
Different testing levels address various aspects of software quality. Unit testing focuses on isolated

Click Here to Buy book on Amazon

modules, while integration testing verifies the interactions between combined components. Alpha testing is performed by users at the developer's site to find major issues, and beta testing occurs at the customer's site in a real-world environment. Each level is essential for ensuring a robust product.

72. (B) White-box testing
Basis path testing is a white-box testing technique that uses the control flow graph of a program to identify a set of linearly independent paths. By executing these paths, testers can ensure that every statement in the code is covered at least once. This method provides a systematic way to achieve high test coverage and identify complex logical errors.

73. (C) Ensure changes have not introduced new faults in existing functionality.
The primary goal of regression testing is to ensure that recent changes or bug fixes have not introduced new faults into existing, previously functioning parts of the software. This involves re-running a suite of test cases after any modification to the code. It is a critical part of the maintenance phase and continuous integration workflows to maintain system stability.

74. (B) The functional requirements and specifications.
Black-box testing focuses on the functional requirements and specifications of the software without any knowledge of its internal code. Testers provide inputs and verify the outputs against the expected behavior defined in the requirements. This approach ensures that the system performs its intended functions correctly from the user's perspective, regardless of how the internal logic is implemented.

75. (C) Performance Testing
Performance testing evaluates how a system behaves under a specific workload, focusing on attributes such as response time, throughput, and resource utilization. It aims to identify bottlenecks and ensure the system can handle its expected traffic efficiently. This type of testing is vital for providing a good user experience and meeting non-functional requirements related to speed and scalability.

76. (B) Stress Testing
Stress testing involves pushing a system beyond its specified operational limits to observe how it behaves under extreme conditions. The goal is to identify the breaking point of the software and ensure it fails gracefully rather than crashing unexpectedly. This type of testing helps developers understand the system's robustness and capacity for recovery in high-pressure scenarios or environments.

77. (A) Both A and R are true and R is the correct explanation of A.
Integration testing is performed after unit testing to identify defects that arise from the interaction between different software modules. While individual components may work correctly in isolation, they might fail when combined due to interface errors or inconsistent data exchange. This phase ensures that the various parts of the system work together seamlessly to fulfill the overall requirements.

NODIA

78. (C) Condition Coverage and Loop Testing
Control structure testing is a category of white-box testing that includes techniques like condition coverage, data flow testing, and loop testing. These methods focus on verifying the logical decisions and repetitive structures within the code. By systematically exercising these control paths, developers can uncover complex errors related to program logic, ensuring the software handles different execution scenarios reliably.

79. (A) Both Statement I and Statement II are true.
Alpha testing is an initial form of acceptance testing conducted in-house by the development or QA team to identify major issues. Once the software is stable, beta testing involves releasing it to a select group of external users for feedback in a real-world setting. Both phases are critical for validating the software against user needs before its final commercial release.

80. (A) To manage and control changes to software artifacts throughout the lifecycle.
Software Configuration Management aims to manage and control changes to software artifacts throughout the entire development lifecycle. By tracking versions, managing baselines, and controlling modifications, SCM ensures the integrity and consistency of the software

to maintain by removing redundancies and improving logic. By regularly refactoring, development teams can prevent technical debt and ensure the long-term health and scalability of their software systems during the development lifecycle.

99. (D) Tracking risk indicators and triggers to see if a risk is becoming likely.
Risk monitoring involves tracking identified risks throughout the project to see if they are becoming more likely or if new risks have emerged. This activity includes watching for triggers and indicators that a risk might manifest. By continuously observing the project environment, managers can implement mitigation plans at the right time, minimizing the potential negative impact on the schedule and budget.

100. (B) P-3, Q-4, R-2, S-1
Fundamental design principles like separation of concerns and functional independence guide the creation of robust software. Separation of concerns breaks complex problems into smaller parts, while functional independence focuses on high cohesion and low coupling. Modularity organizes the system into self-contained units, and information hiding protects implementation details. Together, these principles ensure the resulting system is maintainable and scalable.

101. (B) Non-Functional (Usability)
Requirements that describe qualities like being intuitive or easy to use are classified as non-functional usability requirements. They do not define specific functions but rather the quality of the user's interaction with the system. These attributes are crucial for user adoption and satisfaction, as they directly impact how effectively and efficiently end users can perform their tasks using the software.

102. (B) A flow graph of the program's control structure.
The primary input for basis path testing is a flow graph that represents the control structure of a program. This graph allows testers to calculate cyclomatic complexity and identify a set of linearly independent paths. By exercising these paths, they can ensure that every statement in the code is executed at least once, providing a systematic approach to structural testing.

103. (A) Both A and R are true and R is the correct explanation of A.
Version control is a critical component of Software Configuration Management because it manages different versions of configuration items. It allows teams to track changes, collaborate without conflict, and revert to previous states if a new modification causes errors. By providing a historical record of the system's evolution, version control ensures project integrity and simplifies the process of managing software updates.

104. (B) II and IV
Function Points measure the functionality delivered to the user and are independent of the programming language, making them useful early in the project. However, FP analysis can be more subjective than simply counting Lines of Code. While LOC is easy to measure after coding, it cannot be accurately predicted during requirements analysis and varies significantly based on coding style.

105. (B) Coupling
Coupling is defined as a measure of the degree of interdependence between two software modules. In high-quality design, the goal is to achieve low coupling, meaning modules are relatively independent. This makes the system easier to understand, test, and maintain, as changes in one module are less likely to impact others. High coupling often indicates a poorly structured and rigid system.

106. (B) An ordered list of everything that is known to be needed in the product.
In Scrum, the product backlog is an ordered list of every feature, function, requirement, enhancement, and fix that is known to be needed in the product. It serves as the single source of requirements for any changes to be made. The Product Owner is responsible for managing the backlog, ensuring that the most valuable items are prioritized for the team.

107. (B) The project's size and criticality.
The Crystal family of agile methodologies emphasizes that project practices should be tailored based on the project's size and criticality. Different colors in the Crystal family represent different levels of ceremony and documentation required for various team sizes. This adaptive

Downloaded from www.pdf.tube

Team is fundamental to the Scrum framework. The Product Owner manages the backlog, the Scrum Master facilitates the process, and the Development Team builds the product. This collaborative arrangement promotes self-organization and accountability, enabling the team to deliver high-quality increments of software in every sprint.

90. (B) Pair Programming
Pair programming is an Extreme Programming practice where two developers work together at a single workstation. One person, the driver, writes the code while the other, the navigator, reviews it in real-time. This technique improves code quality, facilitates knowledge sharing, and reduces errors. It is a highly collaborative approach that fosters technical excellence and team cohesion during development.

91. (A) I, III, II, IV
Software testing typically follows a logical progression from individual components to the entire system. It begins with unit testing, followed by integration testing to check module interactions. Next, system testing evaluates the complete product against its requirements. Finally, acceptance testing is performed to ensure the software meets user needs and business objectives. This staged approach helps identify defects early.

92. (C) A human mistake (error) leads to a defect (fault) causing a failure.
The relationship between an error, a fault, and a failure follows a logical sequence. A human mistake during development is an error, which creates a defect or fault in the code. When this faulty code is executed, it can lead to a failure, which is the observable incorrect behavior of the system. This chain of events defines how bugs manifest.

93. (C) Agile Scrum Model
The Agile Scrum model is ideal for small teams working with rapidly changing requirements and requiring frequent delivery of working software. Its iterative nature allows teams to pivot quickly based on feedback, while the sprint structure ensures a steady pace of delivery. This flexibility makes Scrum more effective than rigid, linear models for projects in dynamic and uncertain environments.

94. (B) A description of interactions between an actor and the system.
The primary output of developing use cases is a detailed description of the interactions between an external actor and the system. These scenarios clarify the system's functional requirements by illustrating how users will achieve their goals. By focusing on behavior rather than internal logic, use cases help ensure the system design aligns with user expectations and business processes.

NODIA

95. (C) Encapsulation
Encapsulation is the object-oriented design principle of bundling data and the methods that operate on that data into a single unit, known as a class. This concept promotes information hiding by restricting direct access to the internal state of an object. By exposing only a controlled interface, encapsulation enhances system modularity and makes the code more robust and maintainable.

96. (C) Reusability
In McCall's quality model, reusability refers to the extent to which a program or its components can be reused in other applications. High reusability reduces development time and costs across multiple projects by leveraging proven assets. This factor encourages developers to create modular, well-documented components that are independent of specific application logic, facilitating their integration into future software systems.

97. (B) It is a formally accepted configuration version serving as a reference.
A baseline is a formally accepted version of a configuration item that serves as a fixed reference point for future development. Once a baseline is established, changes to it must follow a formal change control process. This ensures that everyone on the team is working with the same version of the system and provides a stable foundation for testing and release.

98. (B) Refactoring
Refactoring is the process of improving the internal structure of existing code without changing its external behavior. It aims to make the code cleaner, more efficient, and easier

It encompasses attributes like ease of use and user satisfaction. A usable system minimizes the training needed for new users and allows experienced users to perform their tasks efficiently, which is critical for the product's success.

44. (A) I, II, and IV
Quality Assurance involves proactive activities that focus on ensuring the processes used to develop products are effective. Quality Control, on the other hand, is a reactive process centered on identifying defects in the actual products through activities like testing. While QA aims to prevent errors from occurring, QC ensures that any bugs that do appear are caught before release.

45. (D) Affordability
The ISO 9126 standard defines six main quality characteristics: functionality, reliability, usability, efficiency, maintainability, and portability. These factors provide a framework for evaluating software quality from both a technical and user perspective. Affordability is not a characteristic defined by this technical standard, as it relates to economic factors rather than the engineering properties of the software itself.

46. (B) To identify, analyze, and control potential problems impacting the project.
Risk management focuses on identifying, analyzing, and controlling potential problems that could negatively impact a software project. This continuous process involves assessing the probability and impact of various risks, ranging from technical issues to resource constraints. By proactively managing these threats, project managers can minimize their effect on the schedule and budget, increasing the likelihood of a successful delivery.

47. (D) Risk Contingency Planning
Risk contingency planning involves developing specific strategies to be implemented if a foreseen risk actually occurs. This proactive approach ensures that the project team is prepared to respond effectively, minimizing disruption to the development process. By having a pre-defined plan in place, organizations can manage the impact of realized risks more efficiently and maintain progress toward their project goals.

48. (A) P-4, Q-3, R-2, S-1
ISO 9126 quality factors address various aspects of software performance and maintenance. Portability measures how easily software moves between environments, while efficiency focuses on resource usage. Maintainability reflects the ease of making modifications, and reliability concerns the software's ability to perform under specific conditions over time. These metrics help developers and stakeholders evaluate the overall technical quality of a system.

49. (D) The probability of failure-free operation in an environment.
Software reliability is defined as the probability of failure-free operation of a program in a specified environment for a given period of time. It is a critical quality factor, especially for systems where failure could have significant consequences. Reliability is often measured using metrics like mean time between failures and is influenced by the quality of the development and testing processes.

50. (A) Risk Mitigation, Monitoring, and Management
The RMMM plan stands for Risk Mitigation, Monitoring, and Management. It is a comprehensive document that outlines the strategies a project team will use to address identified risks. Mitigation focuses on prevention, monitoring tracks the status of risks, and management defines the actions to take if a risk becomes a reality. This plan is essential for effective project governance.

51. (A) Both Statement I and Statement II are true.
Quality Assurance is a proactive process that aims to prevent defects by focusing on the development methods and standards. In contrast, Quality Control is a reactive process dedicated to finding and fixing defects in the final product through testing and inspections. Both are necessary for high-quality software, as they address different aspects of the development lifecycle to ensure excellence.

52. (B) Portability
In McCall's quality model, portability is the effort required to transfer a program from one hardware or software environment to another. This factor is increasingly important in today's

Click Here to Buy book on Amazon

approach ensures that the development process remains efficient without being overly bureaucratic, matching the specific needs and risks of each project.

108. (C) Faster delivery and increased productivity
A primary benefit of the component-based development model is faster delivery and increased productivity. By using pre-built and tested components, teams can assemble systems more quickly than if they were writing all code from scratch. This approach also promotes software reuse and can lead to higher reliability, as components are often used and verified across multiple different projects and environments.

109. (C) Learning from mistakes and successes is a fundamental part of the process.
The speculate, collaborate, and learn cycle in Adaptive Software Development emphasizes learning from both successes and failures. Speculation replaces traditional planning to acknowledge uncertainty, while collaboration focuses on communication among stakeholders. Learning is the fundamental stage where the team reflects on the results of their work and adapts their approach. This iterative process is essential for managing projects with high levels of complexity.

110. (D) Spiral Model
The spiral model is particularly well-suited for large, complex, and high-risk projects because it incorporates explicit risk analysis in each iteration. It combines the iterative nature of prototyping with the controlled aspects of the waterfall model. By addressing risks early and often, the spiral model helps prevent project failure and ensures that the final product meets its objectives and quality standards.

111. (D) I, II, III, IV, and V
Feature Driven Development is characterized by five main steps: developing an overall model, building a feature list, planning by feature, designing by feature, and building by feature. This model organizes work into small, client-valued pieces of functionality. By focusing on features, FDD provides a clear structure for reporting progress and ensures that the development efforts are directly aligned with user needs.

112. (B) Verification
Verification is the activity of checking whether the software conforms to its technical specifications and requirements. It ensures that the product is being built correctly according to the design and documentation. This process includes activities like inspections, reviews, and walkthroughs. Verification differs from validation, which focuses on whether the software meets the actual needs and expectations of the end users.

113. (C) Modularity facilitates parallel development and simplifies testing.
Modularity involves breaking a system into independent units, which facilitates parallel development by allowing different teams to work on separate modules simultaneously. It also simplifies testing and maintenance, as each module can be verified in isolation. Good design focuses on high cohesion and low coupling within these modules to ensure they remain functionally independent and easy to manage throughout the software lifecycle.

114. (C) Temporal Cohesion
Temporal cohesion occurs when the elements of a module are grouped together because they are all executed within the same time span. An example is a module that performs all system initialization tasks. While these tasks may not be functionally related, their timing creates a logical grouping. However, temporal cohesion is considered weaker than functional cohesion because the internal elements lack a strong relationship.

115. (C) Formal Methods
For safety-critical systems like air traffic control, where correctness is paramount, formal methods are the most appropriate approach. These methods use mathematical specifications and proofs to ensure the system behaves exactly as intended without errors. By providing a rigorous foundation for design and verification, formal methods minimize the risk of failure in environments where human life or significant resources are at stake.

116. (B) A point in the timeline marking completion of a major deliverable.
A milestone is a significant point in a project timeline that marks the completion of a

Downloaded from www.pdf.tube

134. (C) Portability
According to the ISO 9126 standard, portability is the ease with which a software system can be adapted to different specified environments. This includes factors like installability, replaceability, and conformance to standards. A portable system can move between different hardware platforms or operating systems with minimal effort, making it more versatile and reducing the long-term costs of maintenance and platform migration.

135. (B) To assign people with the right skills to specific project tasks.
Staffing in project scheduling involves assigning individuals with the appropriate skills and experience to specific project tasks. This ensures that the work is performed by qualified personnel, which is essential for meeting quality and timeline goals. Effective staffing also considers team availability and workload, preventing burnout and ensuring that the project remains on track throughout its various development phases and milestones.

136. (C) Unit Testing
Unit testing involves the isolated testing of individual software components, such as sub-programs, procedures, or classes. The goal is to verify that each small piece of code functions correctly according to its specific requirements. By catching errors at this level, developers can prevent them from propagating to higher stages of integration and system testing, where they become much harder to identify.

137. (B) Git
Git is a widely used tool for version control in software configuration management. It allows developers to track changes to source code over time, collaborate on the same project using branches, and merge their work seamlessly. By providing a decentralized system for managing file versions, Git helps teams maintain project integrity, prevent data loss, and coordinate their efforts effectively during development.

138. (B) Prototyping Model
The prototyping model is most effective when requirements, especially for the user interface, are unclear or poorly understood. By building a simplified, working version of the system, developers can gather immediate feedback from users and clarify their needs. This iterative process helps ensure that the final product meets user expectations and reduces the risk of building a system that does not satisfy requirements.

139. (C) Anyone or anything outside the system that interacts with the system.
In the context of use cases, an actor is anyone or anything external to the system that interacts with it to achieve a goal. This includes human users, as well as external software systems or hardware devices. Identifying all actors is essential for defining the system's boundaries and ensuring that all possible interactions and requirements are captured during the elicitation process.

NODIA

140. (C) Polymorphism
Polymorphism is an object-oriented design concept that allows different objects to respond to the same message or method call in their own unique ways. This enables a single interface to represent different underlying forms or types of data. Polymorphism increases the flexibility and extensibility of software systems by allowing new classes to be added with minimal changes to existing code structures.

141. (B) Quality Assurance
When a quality assurance audit finds that the team is not following established coding standards, it represents a failure in the Quality Assurance process. QA is responsible for defining and ensuring adherence to the processes and standards used during development to prevent defects. If these processes are ignored, the overall quality of the product is compromised, leading to increased technical debt and errors.

142. (C) Risk Transfer
Purchasing insurance to cover financial losses from a data breach is an example of risk transfer. In this strategy, the potential impact of a risk is shifted to a third party, such as an insurance company. While the risk still exists, the financial burden is shared or removed from the project, allowing the organization to manage its exposure more effectively.

143. (A) Both Statement I and Statement II are true.
White-box testing requires the tester to have detailed knowledge of the program's internal logic and code structure to design effective test

A key activity in Quality Assurance is defining the software development process and establishing quality standards for the team to follow. This proactive approach aims to prevent defects from occurring in the first place. In contrast, Quality Control focuses on identifying existing defects through testing and reviews. QA ensures that the right processes are in place to deliver a high-quality product consistently.

126. (A) Both A and R are true and R is the correct explanation of A.
Risk management must be a continuous process throughout the project lifecycle because the project environment is dynamic. New risks can emerge as development progresses, and the probability or impact of existing risks can change. By regularly identifying and reassessing threats, project managers can adapt their mitigation strategies and ensure that potential problems are addressed before they can significantly derail the project.

127. (B) COCOMO Model
Estimation models that use a formula involving the software size (like KLOC) raised to a constant power are typically variants of the COCOMO model. This algorithmic approach calculates the effort required for a project based on historical data and various cost drivers. COCOMO is widely used because it provides a structured and repeatable way to estimate resources for software development projects.

128. (C) Alpha testing is done at the developer's site, while beta testing is done at the customer's site.
The key difference between alpha and beta testing is the environment in which they occur. Alpha testing is conducted in-house at the developer's site, often using a controlled environment to find major bugs. Beta testing, however, is performed by actual users at their own sites under real-world conditions. This allows the team to gather feedback on performance and usability before release.

129. (B) Re-engineering
Software re-engineering involves modifying an existing system to improve its quality attributes, such as maintainability and reliability, without changing its core functionality. This process often includes reverse engineering to understand the system and forward engineering to implement improvements. It is a cost-effective way to modernize legacy systems, making them easier to support and extend in response to changing business needs.

130. (C) The software is designed with modules having high cohesion and low coupling.
Functional independence in software design is achieved by creating modules with high cohesion and low coupling. High cohesion ensures each module is focused on a single task, while low coupling minimizes dependencies between modules. This independence makes the software easier to understand, develop, and maintain, as changes to one module are unlikely to cause errors in others, improving overall system robustness.

131. (A) Actions, tasks, and work products
In a generic software process model, a framework activity is a high-level stage composed of related actions, tasks, and work products. These elements define the actual work to be done and the results to be produced during that phase. Following this structured approach ensures that every part of the software development lifecycle is clearly defined, measurable, and aligned with project goals.

132. (C) Design constraint (a type of non-functional requirement)
A requirement that specifies a particular technology or standard, such as using AES-256 encryption, is classified as a design constraint. This is a type of non-functional requirement that limits the options available to the designers. Constraints are essential for ensuring that the system integrates with existing infrastructure or meets specific security and regulatory standards required by the organization or client.

133. (A) I, II, and IV
Core design principles for high-quality software include achieving high cohesion and low coupling. High cohesion ensures elements within a module work together toward a single goal, while low coupling reduces dependencies between different modules. Information hiding further improves design by concealing internal implementation details. These concepts work together to create systems that are modular, maintainable, and resilient to changes over time.

Click Here to Buy book on Amazon

across different platforms. This makes it particularly useful early in the software lifecycle, as it can be calculated based on requirements before any code is written, providing more reliable data for planning.

170. (A) Inheritance
Inheritance is an object-oriented design concept that allows a new class, known as a subclass, to inherit properties and behaviors from an existing class, called a superclass. This promotes code reuse and establishes a hierarchical relationship between classes. By allowing developers to build upon existing code, inheritance simplifies the creation of specialized objects and enhances the overall modularity and maintainability of the software system.

171. (C) Synchronize activities and create a plan for the next 24 hours.
The daily Scrum is a short, time-boxed meeting for the development team to synchronize their activities and plan their work for the next twenty-four hours. Each member shares their progress, what they plan to do next, and any impediments they are facing. This meeting fosters transparency, promotes collaboration, and allows the team to identify and address issues quickly, ensuring they remain on track.

172. (C) Unit Testing - Tests interactions between modules.
Unit testing is the process of testing individual software modules or components in isolation to verify they work correctly. It is not about testing the interactions between modules, which is the focus of integration testing. White-box and black-box techniques are used to ensure that internal logic and external functions are correct. Regression testing ensures that changes do not introduce new faults into existing features.

173. (C) Risk Management
The RMMM plan is a key document produced during the risk management process. It outlines the strategies for risk mitigation, monitoring, and management for a specific software project. By detailing how the team will address potential threats, the RMMM plan provides a proactive framework for minimizing project disruptions. It is an essential tool for project managers to ensure a smooth and successful development cycle.

174. (D) Waterfall Model
The waterfall model is a classic prescriptive process model characterized by a linear, sequential approach to software development. It defines a set of distinct phases that must be completed in order. This model prioritizes structure, discipline, and documentation. In contrast, agile models like Scrum, XP, and Crystal focus on iterative delivery and flexibility, making them more adaptive to changing requirements than traditional prescriptive models.

175. (A) Low Coupling
Low coupling is the design principle that advocates for minimizing the connections and dependencies between different software modules. By keeping modules independent, developers can ensure that changes to one part of the system have a minimal impact on others. This modularity simplifies testing, enhances reusability, and makes the overall system more stable and easier to maintain throughout its lifecycle, leading to higher quality software.

176. (B) Boundary Value Analysis
Boundary value analysis is a black-box testing technique that focuses on testing the edges of the input domain. For a password field between eight and sixteen characters, testing lengths of seven, eight, sixteen, and seventeen targets the boundaries where errors are most likely to occur. This method is highly effective for identifying defects related to range checks and off-by-one errors in the program's logic.

177. (C) Statement I is false, but Statement II is true.
Reverse engineering is the process of analyzing a system to understand its structure and logic, while software re-engineering involves modifying a system to improve its quality attributes like maintainability or performance. Reverse engineering is often a step within the larger re-engineering process. Both are important for managing legacy systems, but they serve different primary goals in the lifecycle of software maintenance and evolution.

178. (B) Gantt Chart
A Gantt chart is a project management tool that uses a horizontal timeline to display project tasks, their durations, and their start and finish dates. It provides a visual representation of the project

cases. In contrast, black-box testing is performed without any knowledge of the internal workings, focusing entirely on functional requirements. Both techniques are necessary for comprehensive software testing, as they identify different types of defects and ensure high overall quality.

144. (C) Reverse engineering
Re-creating design and specification documents from existing source code is a core objective of reverse engineering. This process is particularly valuable for legacy systems that have little to no up-to-date documentation. By analyzing the code, developers can understand the system's architecture and logic, providing the information needed for maintenance, re-engineering, or migrating the software to a newer technology platform or environment.

145. (B) An activity is comprised of actions, and each action is defined by a Task Set.
In the generic process model, a framework activity is a high-level stage composed of multiple actions. Each of these actions is defined by a specific task set, which outlines the actual work, work products, and quality assurance points required. This hierarchical structure ensures that the high-level goals of each framework activity are achieved through the systematic execution of well-defined technical tasks.

146. (C) The system shall allow a user to add items to a shopping cart.
A requirement stating that the system must allow a user to add items to a shopping cart is a classic functional requirement. It specifies a specific service or feature that the software must provide to fulfill user needs. Functional requirements define what the system does, while non-functional requirements describe quality attributes such as performance, security, availability, and usability of those functions.

147. (B) Communicational Cohesion
Communicational cohesion occurs when the parts of a module are grouped together because they all operate on the same data structure or information. This form of cohesion is stronger than temporal or logical cohesion because the elements share a common data-related goal. Modules with communicational cohesion are relatively easy to understand and maintain, as their internal logic is centered around a specific piece of information.

148. (B) The discussion between team and owner to clarify user story details.
In the 3 Cs of a user story, the Conversation refers to the ongoing discussion between the development team and the product owner to clarify details. These interactions help ensure everyone has a shared understanding of the requirement before work begins. This collaborative approach replaces lengthy, static documentation with dynamic communication, allowing for better alignment and responsiveness to changing needs during the sprint.

149. (C) Product Operation
In McCall's quality model, correctness and reliability are categorized as factors related to product operation. These factors measure how well the software performs its functions and how consistently it does so during its actual use. Product operation also includes attributes like efficiency, integrity, and usability, which together define the quality of the software from the perspective of the end user during execution.

150. (A) P-3, Q-2, R-4, S-1
Different types of risk impact software projects in various ways. Project risks affect the schedule and resources, while product risks concern the quality and functionality of the software. Business risks relate to market viability and competition, and technical risks involve challenges with technology and tools. Identifying and managing all these risk categories is essential for ensuring the overall success and delivery of the project.

151. (B) Branch/Decision Coverage
Branch or decision coverage is a testing criterion that requires every possible outcome of each decision point in the code to be executed at least once. This means both true and false paths for every if statement or case must be tested. Achieving high branch coverage ensures that the software's logical structure is thoroughly exercised, helping to identify errors in decision-making logic.

152. (C) Software Configuration Management
Managing a baseline and controlling changes to it is a central function of Software

Downloaded from www.pdf.tube

Configuration Management. A baseline provides a stable reference point for development, and any modifications must be formally reviewed and approved through a change control process. This ensures that the system's integrity is maintained and that all team members are working with consistent versions of the software and its documentation.

153. (C) A meeting to inspect performance and create a plan for improvement.
A sprint retrospective is a Scrum meeting held at the end of a sprint to inspect the team's performance and identify opportunities for improvement. The team discusses what went well and what could be improved in terms of processes, tools, and relationships. This meeting is crucial for the Inspect and Adapt pillar of Scrum, fostering continuous improvement and team growth over time.

154. (C) It is rigid and does not easily accommodate requirement changes.
The primary drawback of the waterfall model is its rigidity and inability to easily accommodate changes in requirements once the project has started. Its sequential nature means that a mistake made in the early stages can be very costly to fix later. This lack of flexibility makes it less suitable for modern projects where user needs often evolve as the software is being developed.

155. (C) Following a plan over responding to change
The Agile Manifesto values responding to change over following a rigid plan. While planning is still important, agile methodologies prioritize flexibility and the ability to adapt as more is learned about the project. Other principles include valuing individuals and interactions, working software, and customer collaboration. By being responsive to change, agile teams can ensure the final product meets the users' evolving needs.

156. (C) High coupling
When one software module directly accesses and modifies the internal data of another, it is a clear example of high coupling. This creates a strong dependency between the two modules, making the system fragile and difficult to maintain. Changes to one module are likely to break the other, increasing the complexity of debugging and reducing the overall modularity and reusability of the system.

157. (A) I and II only
Key goals of user interface design include making the interface aesthetically pleasing and, more importantly, efficient and easy to use. A well-designed UI reduces user errors and enhances satisfaction by focusing on the user's workflow. While technical factors like cross-device compatibility are important, the core mission of UI design is to facilitate effective and effortless interaction between the user and the software.

158. (A) I, II, and IV only
McCall's quality model includes factors such as portability, reliability, and usability, which assess software quality from different viewpoints. These factors are grouped into categories like product operation, revision, and transition. Economic factors like cost and project management concerns like development time are not considered technical quality factors within this specific model, which focuses on the engineering characteristics of the software itself.

159. (A) Organic
The organic mode of the COCOMO model is used for projects with small, highly skilled teams and flexible requirements. These projects are typically straightforward and involve familiar technology. Because the team is experienced and the problem is well-understood, the development process requires less formal overhead than more complex projects. This mode results in lower effort and cost estimates compared to the semi-detached or embedded modes.

160. (A) Both A and R are true and R is the correct explanation of A.
Regression testing is essential during software maintenance to ensure that modifications like bug fixes or enhancements do not negatively impact existing functionality. By re-running tests, developers can catch unintended side effects caused by code changes. This process maintains system stability and prevents the introduction of new defects into previously working parts of the software, which is critical for long-term reliability and satisfaction.

161. (B) Forward Engineering

Forward engineering is the traditional process of creating a software system from high-level requirements through analysis, design, and coding to an executable product. It involves moving from more abstract representations to a concrete implementation. This is the standard path for building new software systems, distinct from reverse engineering, which extracts design information from code, or re-engineering, which modernizes existing software systems.

162. (A) The progression of a software process through its various stages.
A process lifecycle refers to the various stages a software process progresses through, from its initial definition to its continuous improvement. It provides a structured framework for managing development activities and ensures that processes are consistently applied across projects. By understanding the lifecycle, organizations can refine their methods, improve efficiency, and ensure that their software engineering practices remain effective and aligned with industry standards.

163. (B) Create models that refine and clarify the elicited requirements.
The primary goal of requirement analysis and modeling is to create abstract models that refine and clarify the requirements gathered during elicitation. These models help stakeholders and developers identify inconsistencies, omissions, and ambiguities before design begins. By providing a clear and shared understanding of the system's intended behavior, modeling ensures that the final software product correctly addresses the stakeholders' actual needs and expectations.

164. (B) A problem should be decomposed into distinct features with minimal overlap.
Separation of concerns in software design implies decomposing a complex system into distinct features or modules with minimal overlap. Each module should handle a specific concern, making the system easier to understand, develop, and maintain. This principle promotes modularity and reduces complexity by isolating different aspects of the software's functionality, which in turn simplifies debugging and enhances the overall robustness of the application.

165. (C) I, II, and III only
Core principles of user interface design include placing the user in control, reducing the user's memory load, and maintaining consistency throughout the interface. Using complex technical jargon is a negative practice that should be avoided. A good UI design focuses on clarity and simplicity, ensuring that users can perform their tasks effectively without being overwhelmed by complicated or unfamiliar terminology or interface elements.

166. (B) Implement a mitigation or avoidance strategy immediately.
When a risk is identified with both high probability and high impact, it poses a significant threat to the project's success. In such cases, the team must immediately implement a mitigation or avoidance strategy to reduce the threat. Ignoring such risks or accepting them without a plan is dangerous, as they are likely to manifest and cause major delays, cost overruns, or failure.

167. (C) II, IV, I, III
The software re-engineering process typically begins with reverse engineering to understand the current system. This is followed by data restructuring and code restructuring to improve the system's internal quality. Finally, forward engineering is used to implement new features or modernize the application. This systematic approach ensures that the revamped system is based on a solid understanding of the original while benefiting from improved design.

168. (C) Beta Testing
Beta testing involves releasing a nearly complete version of the software to a select group of external users for testing under real-world conditions. This allows the development team to gather feedback on usability, performance, and stability before the final release. By testing in various environments and with different user behaviors, beta testing helps identify elusive bugs that may not have appeared during internal testing.

169. (A) It is language-independent and technology-independent.
The most significant advantage of Function Point analysis over Lines of Code is its independence from technology and programming languages. FP measures functionality from the user's perspective, allowing for consistent estimation

Click Here to Buy book on Amazon

major deliverable or a key stage of the project. Milestones are used to track progress and provide a high-level view of the schedule to stakeholders. Unlike individual tasks, milestones have zero duration and represent a successful outcome or the transition between different project phases.

117. (B) III, II, I
COCOMO classifies projects into three modes based on size and complexity: organic, semi-detached, and embedded. Organic projects are small and simple, while semi-detached projects have intermediate complexity. Embedded projects are the most complex, often involving rigid hardware and software requirements. Arranging them in increasing order of difficulty results in organic, then semi-detached, and finally the embedded mode of development.

118. (D) Statement Coverage
Black-box testing techniques like equivalence partitioning, boundary value analysis, and decision table testing focus on the system's functional specifications. Statement coverage, however, is a white-box testing metric that measures the proportion of source code executed during testing. Because statement coverage requires knowledge of the internal code structure to calculate, it is not categorized as a black-box or functional testing technique.

119. (B) Reverse Engineering
Analyzing a legacy system's source code to understand its functionality and create documentation is a primary application of reverse engineering. This process helps developers gain insights into systems that may have lost their original design documents or expertise over time. By reconstructing a higher-level view of the software, reverse engineering facilitates future maintenance, migration, or re-engineering efforts for older applications.

120. (C) P-3, Q-2, R-1, S-4
Different software process models are suited for specific project types. The waterfall model works best for stable requirements, while prototyping helps clarify unstable needs. The spiral model is ideal for high-risk systems due to its iterative risk analysis. Agile Scrum is designed for rapid delivery in changing environments. Matching the model to the project's characteristics is essential for ensuring a successful development outcome.

121. (C) Aspect-Oriented Software Development
Aspect-Oriented Software Development is designed to address cross-cutting concerns that affect multiple modules, such as security, logging, and transaction management. These concerns often tangle the core logic in traditional object-oriented designs. AOSD allows developers to modularize these aspects separately, improving code clarity and maintainability. This paradigm provides a more effective way to manage systemic properties that do not fit neatly into classes.

122. (C) Statement I is false, but Statement II is true.
The Software Requirements Specification is a formal document shared with the customer to define the project's scope and requirements. It is not an internal-only document. Meanwhile, a use case provides a detailed description of how a user interacts with the system to achieve a specific goal. Use cases are valuable for capturing functional requirements and ensuring the system meets user needs.

123. (A) I, III, and IV only
Essential characteristics for a good Software Requirements Specification include correctness, completeness, and verifiability. These attributes ensure that the document accurately reflects stakeholder needs and can be used as a basis for testing. Ambiguity and high volatility are negative traits that can lead to project failure. A high-quality SRS provides a clear, stable foundation for all subsequent design and development activities.

124. (C) Designing the data structures and database schema to be used.
Data design is a phase of software design that focuses on defining the data structures and database schema required by the system. It involves modeling how information will be stored, organized, and accessed by various software components. Effective data design ensures that the system can handle its data requirements efficiently and provides a solid foundation for the implementation of application logic.

125. (C) Defining the software development process and standards.

NODIA

schedule, allowing managers to track progress and identify dependencies between activities. By showing how terminal and summary elements relate over time, Gantt charts help teams stay organized and meet deadlines.

179. (C) Usability
In the ISO 9126 quality standard, usability refers to the software's capability to be understood, learned, used, and perceived as attractive by the user. It focuses on the quality of the user experience and the ease with which users can achieve their goals. A usable system minimizes the learning curve and provides an efficient and satisfying interface, which is critical for the product's overall success.

180. (A) I, III, and V only
Agile development is characterized as an iterative and incremental approach that values face-to-face communication and frequent delivery of working software. Unlike traditional models, agile methodologies welcome changes in requirements throughout the project. They prioritize individuals and interactions over rigid processes and comprehensive documentation. This flexibility allows teams to respond effectively to user feedback and deliver products that truly meet the stakeholders' needs.

181. (C) Representing stakeholders and managing the backlog to maximize value.
The Product Owner in Scrum represents the interests of the stakeholders and is responsible for managing the product backlog. Their primary goal is to maximize the value of the product by prioritizing the most important features for the development team. They act as a bridge between the business side and the technical team, ensuring that the software being built aligns with organizational objectives.

182. (B) Temporal Cohesion
Temporal cohesion occurs when tasks are grouped within a module simply because they must be executed at the same time, such as during system initialization or shutdown. Although these tasks may not be functionally related, their timing creates a logical association. Temporal cohesion is considered one of the weaker forms of cohesion because the internal elements lack a strong functional or logical relationship.

183. (A) Both A and R are true and R is the correct explanation of A.
A good software design exhibits high cohesion because it means the elements within a module are strongly related and focused on a single task. This clarity makes the module easier to understand, test, and maintain. Furthermore, highly cohesive modules are more likely to be reusable in other parts of the system or in different projects, which improves the overall efficiency of software development.

184. (B) Product-oriented inspection and testing to find defects.
Quality Control activities focus on product-oriented inspection and testing to identify defects in the final software. Unlike Quality Assurance, which focuses on the development process, QC is dedicated to finding bugs in the code before it is released to the customer. This reactive approach ensures that the delivered product meets the specified quality standards and performs its intended functions correctly and reliably.

185. (B) Estimation by Analogy
Estimation by analogy is a software sizing technique that uses historical data from past projects to estimate the size of a new, similar project. By comparing the characteristics and complexity of previous work to the current requirements, managers can derive more accurate effort and cost estimates. This method relies on the team's experience and the availability of well-documented records of past performance.

186. (B) P-3, Q-4, R-1, S-2
Software testing is organized into levels, each with a specific focus. Unit testing verifies individual functions, while integration testing checks the interfaces between modules. System testing evaluates the entire product against its technical requirements. Finally, acceptance testing validates the software against the actual business needs and user expectations. Each level is vital for ensuring that the final product is robust, functional, and valuable.

187. (D) Change Control
Change control is the systematic process of tracking and managing modifications to software artifacts. It involves evaluating proposed changes, assessing their impact, and formally

NODIA

approving or rejecting them. This ensures that every update is documented and that the project baseline remains stable. By controlling changes, teams can prevent conflicting modifications and maintain the integrity and quality of the software throughout its development lifecycle.

188. (D) Financing
The generic software process model consists of framework activities like communication, planning, modeling, construction, and deployment. These phases represent the technical and management actions required to develop software. Financial activities, such as securing project funding or managing budgets, are business-related functions and are not considered part of the core engineering process phases defined in this technical framework for software development and delivery.

189. (C) A quality attribute or constraint, such as performance or security.
A non-functional requirement describes a quality attribute or constraint of the system, such as performance, security, or reliability. Rather than defining specific features, it specifies how the system must perform its functions. For example, a requirement for the system to process data within a certain time limit is non-functional. These requirements are essential for ensuring the system is efficient and meets user standards.

190. (C) Coincidental Cohesion
Coincidental cohesion is the weakest form of cohesion, occurring when the parts of a module are grouped together with no meaningful relationship. This often happens as a result of poor design or arbitrary modularization. Modules with coincidental cohesion are difficult to understand, maintain, and reuse because they lack a clear purpose. High-quality software design strives to replace coincidental cohesion with stronger forms.

191. (D) Modularity
McCall's quality model identifies several factors like integrity, flexibility, and testability to evaluate software quality. These are grouped into product operation, revision, and transition. Modularity is a design principle rather than a quality factor in this specific model. While modularity contributes to factors like maintainability and flexibility, it is not listed as an independent top-level factor within McCall's established quality framework.

192. (B) Quantifying the product size as an input to effort and cost estimation.
Software sizing is the process of quantifying the magnitude of a software product as an input to effort and cost estimation. Common metrics used for sizing include Lines of Code and Function Points. By determining the size of the system early, project managers can better plan resources, schedule tasks, and provide stakeholders with more realistic expectations for the project's overall development timeline.

193. (B) Guarantee every linearly independent path is executed at least once.
The primary goal of basis path testing is to identify and execute a set of linearly independent paths through the program's control structure. This ensures that every individual statement and every decision outcome is exercised at least once. By achieving this level of coverage, testers can systematically uncover logical errors and ensure the software's structural integrity, which is vital for overall reliability.

194. (B) I, III, and IV only
Requirements elicitation uses various techniques like interviews, use cases, and prototyping to gather information from stakeholders. Interviews provide direct insights, use cases illustrate system behavior, and prototypes offer working models for feedback. Technical tasks like code reviews or unit testing are part of the development and quality control phases, respectively, and are not typically used to discover the initial user requirements or needs.

195. (B) Decomposing a large problem into smaller independent parts.
The separation of concerns principle promotes decomposing a large and complex problem into smaller, independent, and manageable parts. Each part, or concern, addresses a specific aspect of the overall functionality. This approach enhances modularity, simplifies development, and makes the system easier to test and maintain. By isolating concerns, developers can focus on individual modules without being overwhelmed by the entire system's complexity.

Downloaded from www.pdf.tube

196. (A) Both A and R are true and R is the correct explanation of A.
Agile models are fundamentally iterative and incremental, delivering software in small functional pieces throughout the project lifecycle. This approach allows teams to be flexible and adapt to changing requirements or feedback as work progresses. By breaking development into manageable cycles, agile methodologies ensure that the final product remains aligned with user needs and can respond effectively to new challenges or insights.

197. (B) Ensuring development conforms to standards to prevent defects.
The primary goal of Software Quality Assurance is to ensure that the development process conforms to established standards to prevent defects. SQA focuses on proactive activities such as process audits, standards definition, and training. By building quality into the process, organizations can reduce the number of bugs found during testing, improve efficiency, and consistently deliver high-quality software products to their customers and stakeholders.

198. (B) Constructive Cost Model
COCOMO stands for Constructive Cost Model, which is a widely used algorithmic model for estimating software development effort, cost, and schedule. It uses formulas based on historical data to predict project resources. Developed by Barry Boehm, COCOMO provides different levels of detail, from basic to detailed, allowing project managers to refine their estimates as they gain more information during the software lifecycle.

199. (B) An error is a human mistake, while a fault is the defect in the code.
An error is a human mistake made during the development of software, such as an incorrect line of code or a misunderstood requirement. A fault, also known as a defect or a bug, is the manifestation of that error within the software artifact. When the faulty code is executed, it can lead to a failure, which is the system's observable incorrect behavior.

200. (B) P-2, Q-1, R-4, S-3
Software Configuration Management includes several key activities for maintaining system integrity. Version control tracks multiple versions of components, while change control manages proposed modifications to the baseline. Configuration audits verify that the product matches the documented configuration, and status reporting keeps stakeholders informed about the current state of items. Together, these processes ensure the software remains organized and consistent throughout its evolution.

201. (C) Feature Driven Development (FDD)
Feature Driven Development is an agile methodology structured around five core activities: developing an overall model, building a feature list, planning by feature, designing by feature, and building by feature. It emphasizes a feature-centric approach to development, delivering small, client-valued pieces of functionality in short iterations. This model is effective for managing larger projects while maintaining the benefits of agile flexibility and focus.

202. (C) Software Requirements Specification (SRS)
The Software Requirements Specification is the formal document that describes the essential functional and non-functional requirements of a software product. It serves as a contract between the client and the developer, ensuring a shared understanding of the project's scope and goals. A well-written SRS is unambiguous and verifiable, providing the definitive foundation for the subsequent design, coding, and testing phases.

203. (C) Functional Independence
Functional independence is the design characteristic most closely associated with cohesion and coupling. It is achieved by creating modules with high internal cohesion and low external coupling. This independence ensures that each module performs a single, focused task and has minimal dependencies on other modules. The result is a system that is easier to understand, develop, test, and maintain over its entire software lifecycle.

NODIA

204. (D) High quality is only achieved by testing at the end of the cycle.
High quality is not achieved only by testing at the end of the development cycle; it must be built into every phase of the project. Quality Assurance focuses on preventing defects in the process, while Quality Control focuses on identifying

222. (B) Requirements Elicitation
Requirements elicitation is the specific process of gathering requirements from clients, users, and other stakeholders through various techniques. This phase is critical for understanding the problem the software is intended to solve. Effective elicitation involves active listening and probing to uncover hidden needs, ensuring that the resulting requirements specification is complete and accurately reflects the actual goals and expectations of the project sponsors.

223. (C) Logical Cohesion
Logical cohesion occurs when a module performs a set of tasks that are related only because they belong to the same logical category, such as different types of tax calculations. While the tasks are conceptually similar, they may not share data or a common functional goal. Logical cohesion is considered weaker than functional cohesion because the internal elements are not tightly integrated toward a single task.

224. (A) I, III, and V only
Quality Control activities are product-oriented and focus on identifying defects. Examples include code inspections, unit testing, and peer reviews, all of which examine the actual artifacts of development to find errors. Defining standards and auditing processes are Quality Assurance activities, which focus on the development methodology itself. QC ensures that individual pieces of the system are correct and meet the required quality standards.

225. (B) To illustrate the schedule, showing tasks, durations, and dependencies.
A time-line chart, such as a Gantt chart, serves as a visual representation of the project schedule, showing tasks, their durations, and their dependencies. It provides a clear overview of when activities start and finish, helping project managers track progress and identify potential delays. This tool is essential for effective resource allocation and for communicating the project timeline and milestones to all stakeholders.

226. (B) Grey-box testing
Grey-box testing is a methodology where the tester has limited knowledge of the system's internal workings. They may have access to design documents or database schemas, which they use to design more effective test cases than pure black-box testing. This approach combines the functional focus of black-box testing with some of the structural insights of white-box testing, helping to identify complex system defects.

227. (B) Reverse engineering understands the system, while re-engineering alters it.
The main difference between reverse engineering and re-engineering is that the former focuses on understanding an existing system, while the latter involves actually altering it. Reverse engineering extracts design information from source code to create higher-level representations. Software re-engineering uses this information to restructure the code and data, improving quality attributes like maintainability. Reverse engineering is often an essential initial phase of any re-engineering effort.

228. (C) Small Releases
The Extreme Programming practice of small releases involves delivering the system in small, frequent increments to the customer. This approach allows for early and continuous feedback, ensuring the software meets user needs and reduces the risk of long-term project failure. Each release is a functional version of the system that provides real value, allowing the business to benefit from the software throughout development.

229. (B) Software Requirements Specification
SRS stands for Software Requirements Specification. It is a fundamental document in software engineering that formally defines all the requirements for a software system, including functional and non-functional aspects. A well-written SRS acts as a contract between the client and the developer, providing a clear and unambiguous baseline for design, coding, testing, and validation of the final product to ensure it meets expectations.

230. (D) High cohesion and low coupling
A software design that is easy to maintain and extend is characterized by high cohesion and low coupling. High cohesion ensures that the elements within each module are closely related to a single task, making them easier to understand. Low coupling minimizes the dependencies between modules, ensuring that changes to one part

Click Here to Buy book on Amazon

on functional requirements. Grey-box testing involves some knowledge of internal structures, such as database schemas. White-box testing requires the most knowledge, as it involves analyzing and exercising the internal source code and logic. This progression helps ensure thorough coverage from multiple perspectives.

214. (C) Construction
In the generic software process model, construction is considered a framework activity. It encompasses the actions of coding and testing the software to build a functional product. Other activities like communication, planning, and modeling provide the foundation, while construction delivers the actual executable software. Each framework activity is essential for a disciplined approach to software development, leading to high-quality results and successful deployment.

215. (A) I and II
Functional requirements define what the system must do, such as specific features or services. Non-functional requirements describe how well the system performs, focusing on quality attributes like speed or security. A requirement to write software in a specific language is a design constraint, a type of non-functional requirement. Use cases are primarily tools for capturing functional behavior rather than non-functional operational constraints.

216. (C) Object-Oriented Design
Object-oriented design centers on the design of classes, their attributes, and their methods, as well as the relationships between them. This approach uses principles like encapsulation, inheritance, and polymorphism to create modular and reusable software components. By modeling real-world entities as objects, OOD provides a structured way to manage complexity and build systems that are flexible and easy to maintain over time.

217. (C) Maintainability
In the ISO 9126 quality model, the capability of a software product to be modified falls under maintainability. This includes attributes like analyzability, changeability, stability, and testability. A maintainable system is designed to be easily updated to fix bugs, improve performance, or adapt to new requirements. High maintainability is critical for reducing the long-term costs of software ownership and ensuring the system remains useful.

218. (C) Embedded
In the COCOMO model, projects involving complex control systems or real-time software are classified under the embedded mode. These projects often have rigid requirements and significant hardware constraints, making them the most difficult to develop. The embedded mode results in the highest effort and cost estimates due to the technical complexity and the high level of coordination and formal procedures required during development.

219. (C) Regression Testing
Regression testing is performed after a bug fix or code modification to ensure that the change has not introduced new faults in other parts of the system. In this case, finding that a profile page fails after a login fix is a clear regression defect. Regular regression testing is vital for maintaining system stability and ensuring that the software remains reliable as it evolves.

NODIA

220. (B) Both A and R are true but R is not the correct explanation of A.
While change control procedures are essential for maintaining the integrity of the project baseline, they are often perceived as bureaucratic and slow. This is because every proposed modification must be systematically evaluated for its impact on the schedule, budget, and other system components before being approved. This tension between the need for control and the desire for speed is a common challenge in management.

221. (B) The Generic Process Model
The generic software process model is characterized by five framework activities: communication, planning, modeling, construction, and deployment. These activities provide a structured roadmap for development teams, guiding them through every stage from initial stakeholder discussions to final product release. By following this disciplined approach, organizations can ensure consistency, manage complexity, and deliver high-quality software that meets its intended objectives and requirements.

them in the product. Models like McCall's and ISO 9126 help define quality factors. Effective quality management requires a continuous, proactive approach throughout.

205. (C) Risk acceptance or retention
Risk acceptance, or retention, occurs when a project team decides to absorb a potential loss rather than implementing a mitigation strategy. This approach is often used for low-impact risks or when the cost of mitigation exceeds the potential loss. While the team acknowledges the risk exists, they do not take active steps to prevent it, instead focusing their resources on more significant threats.

206. (C) Statement I is false, but Statement II is true.
Basis path testing is a white-box testing technique, not a black-box one, because it requires knowledge of the program's internal logic. It uses the cyclomatic complexity of a control flow graph to determine the number of linearly independent paths through the code. By executing these paths, testers can ensure that every statement is covered at least once, providing a rigorous structural test.

207. (B) To improve maintainability and understandability of a legacy system.
A primary motivation for software re-engineering is to improve the maintainability and understandability of legacy systems. As older software becomes difficult to support due to outdated technology or lack of documentation, re-engineering provides a way to modernize the system without starting from scratch. By restructuring the code and data, developers can extend the life of valuable software assets and reduce maintenance costs.

208. (C) Trusting the quality and reliability of third-party components.
A major challenge in component-based development is trusting the quality and reliability of third-party or commercial off-the-shelf components. Since the source code of these components is often unavailable, developers must rely on the vendor's documentation and testing. This can introduce risks related to security, performance, and long-term support. Careful evaluation and testing of components are essential to ensure they integrate successfully into the system.

209. (C) Web Engineering
While methodologies like Scrum, DSDM, and Adaptive Software Development are recognized agile process models, web engineering is a specific discipline focused on the development of web-based systems. Web engineering may incorporate agile practices, but it is defined by its focus on the unique challenges of the web environment rather than being a general agile methodology itself. It addresses concurrency, security, and scalability.

210. (C) Agile Models (like Scrum and XP)
Agile models like Scrum and Extreme Programming rely heavily on user stories to capture requirements from the end user's perspective. A user story is a simple, natural language description of a feature and the value it provides. This format encourages communication and collaboration between the development team and stakeholders, allowing requirements to be refined and prioritized throughout the project in short iterations.

211. (C) I, II, and III only
Good architectural design provides several benefits, including improved communication among stakeholders, a framework for system analysis, and the facilitation of large-scale software reuse. It offers a high-level view of the system's structure and behavior, ensuring that the software meets its quality requirements. While detailed procedural logic is developed during component-level design, the architecture provides the essential blueprint for the entire development project.

212. (C) Software Reuse
Maintaining a library of tested and certified software components for use across different projects is a direct application of software reuse. This practice increases productivity and ensures consistency by leveraging proven assets. By reusing certified components, organizations can reduce the risk of introducing errors and decrease the time and effort required to develop new systems, leading to more efficient and reliable software production.

213. (C) II, III, I
Testing techniques can be arranged by the degree of internal knowledge required. Black-box testing requires the least, as it focuses entirely

Downloaded from www.pdf.tube

of the system have a minimal impact on other components during development.

231. (B) Efficiency
In the ISO 9126 quality standard, efficiency refers to the software product's performance relative to the amount of resources used, such as CPU time, memory, and network bandwidth. It ensures that the system meets its performance requirements without wasting resources. Improving efficiency is critical for providing a good user experience and for ensuring the software can run effectively on its intended hardware and environment.

232. (B) The project exceeds its budget and schedule.
A project risk is an event or condition that could negatively impact the project's schedule, budget, or resources. An example is a project exceeding its budget and schedule. These risks differ from product risks, which concern the technical quality of the software, and business risks, which relate to market viability. Identifying and managing project risks is a primary responsibility of the project manager.

233. (A) Both Statement I and Statement II are true.
Unit testing is typically a white-box technique where developers test individual code modules with full knowledge of their internal logic. Integration testing, however, can use both black-box techniques to verify functional requirements and white-box techniques to check the technical interfaces and data flow between modules. Using a combination of these approaches ensures that the software is thoroughly tested from both structural and functional perspectives.

234. (B) To create a separate line of development for new features.
Branching in version control allows developers to create separate lines of development from the main codebase. This is useful for working on new features, bug fixes, or experiments without affecting the stable version of the software. Once the work is complete and tested, the branch can be merged back into the main line, ensuring a controlled and organized evolution of the software system.

235. (C) Prescriptive, heavyweight models like Waterfall
The Manifesto for Agile Software Development was created as a direct response to the perceived overhead and rigidity of prescriptive, heavyweight models like the waterfall model. These traditional approaches often emphasized documentation and process over working software and human interactions. Agile methodologies seek to streamline development by prioritizing flexibility, customer collaboration, and the delivery of functional software in short, iterative development cycles.

NODIA

236. (A) P-1, Q-3, R-2, S-4
Understanding the distinctions between testing terms is essential for quality assurance. An error is a human mistake, which leads to a fault, a static defect in the code. A failure occurs when this fault is executed, resulting in observable incorrect behavior. Bug is a common industry term synonymous with a fault. This sequence defines how defects move from a developer's mind to the software product.

237. (B) It uses timeboxing to fix time and resources, adjusting functionality.
A key characteristic of the Dynamic System Development Model is the use of timeboxing to fix project time and resources while allowing functionality to be adjusted. DSDM prioritizes the delivery of the most important business needs within a set timeframe. This agile methodology ensures that projects remain on schedule and within budget, even if it means some lower-priority features are deferred to later releases.

238. (A) Both A and R are true and R is the correct explanation of A.
Information hiding is a fundamental design principle because it reduces module coupling by concealing the internal implementation details from other parts of the system. This creates a clear separation between a module's interface and its logic. Consequently, changes to the internal workings do not affect other components as long as the interface remains stable, making the software more robust and easier to maintain.

239. (B) I, II, and III only
Software project management activities include project scheduling, risk management, and cost estimation. these tasks focus on the planning, organizing, and controlling of resources to

achieve specific goals within the project's constraints. while technical activities like testing and elicitation are essential to the software process, they are typically categorized as engineering tasks rather than the administrative and strategic functions of project management and governance.

240. (C) Subjecting the system to extreme loads to observe failure.
Stress testing involves subjecting a software system to extreme loads or conditions that exceed its specified operational limits. The goal is to observe how the system fails and to ensure it does so gracefully without catastrophic data loss or crashing. This type of testing helps identify the ultimate capacity of the software and provides insights into its robustness and error-handling capabilities.

241. (B) Using standardized, documented, and componentized assets.
Software reuse is most effectively implemented through the use of standardized, documented, and componentized assets. Rather than simple copying and pasting, this disciplined approach involves creating modular components that are designed and verified for use in multiple projects. By leveraging these high-quality assets, organizations can improve productivity, reduce development costs, and ensure a higher level of consistency and reliability in their software systems.

242. (B) Prototyping and Waterfall
The spiral model of software development combines the iterative nature of the prototyping model with the controlled, systematic aspects of the waterfall model. It is organized into a series of iterations, each incorporating risk analysis to identify and address potential problems early. This hybrid approach makes the spiral model particularly effective for large, complex, and high-risk projects where rigorous management and flexibility are required.

243. (C) The system must process a transaction in under 3 seconds.
A requirement that the system must process a transaction in under three seconds is a classic example of a non-functional requirement related to performance. It specifies a quality attribute or constraint on the system's operation rather than a specific function. These requirements are essential for ensuring that the software meets user expectations for speed and responsiveness, which are critical for overall satisfaction.

244. (C) Sequential Cohesion
Sequential cohesion occurs when the elements of a module are grouped because the output of one element serves as the input to the next in a specific order. This creates a strong relationship between the tasks within the module. While not as strong as functional cohesion, sequential cohesion represents a logical and efficient way to organize related processing steps within a software component.

245. (C) Risk Management
The RMMM plan, which stands for Risk Mitigation, Monitoring, and Management, is the primary output of the risk management activity. This document details the strategies the project team will use to address potential threats to the project's success. By providing a structured approach to risk, the RMMM plan helps ensure that the project remains on track and that resources are managed effectively.

246. (B) To visually represent the schedule and track progress against it.
The primary purpose of a time-line chart, such as a Gantt chart, is to visually represent the project schedule and track progress against it. It shows individual tasks, their durations, and their dependencies, providing an overview of the project's timeline. This tool helps project managers identify delays, manage resources efficiently, and communicate the project's status and milestones to stakeholders throughout the development lifecycle.

247. (D) Test-Driven Development (TDD)
Test-Driven Development is a software development practice where tests for a specific feature are written before the feature's code itself. This ensures that the requirements are well-understood and that every piece of new code is verified from the start. TDD leads to cleaner, more modular code and higher test coverage, which significantly improves the overall quality and maintainability of the software system.

Click Here to Buy book on Amazon

248. (A) IV, II, V, I, III
The logical order of generic process framework activities begins with communication to gather requirements. This is followed by planning for resources and schedules, and modeling to design the system. Once these phases are complete, the process moves into construction, where the code is written and tested, and finally deployment, where the software is delivered to the customer for its actual use and feedback.

249. (B) I, II, III, and IV only
Key software design concepts include abstraction, modularity, information hiding, and the measures of cohesion and coupling. These principles guide the creation of well-structured and maintainable systems. Abstraction focuses on essentials, modularity organizes components, and information hiding protects data. Cohesion and coupling evaluate module quality. Together, these concepts provide a technical framework for effective software engineering, while requirement elicitation is a phase that precedes design.

250. (A) Both A and R are true and R is the correct explanation of A.
Both verification and validation are necessary to ensure software quality. Verification ensures the product is built correctly according to technical specifications, while validation ensures the right product is built to meet actual user needs. Together, they provide a comprehensive approach to quality management, addressing both the technical accuracy of the development process and the overall business value and utility of the final product.

251. (B) Framework activities are overarching phases, and each is composed of specific task sets
Framework activities within the Generic Process Model represent overarching phases that define the development lifecycle. Each of these high-level activities is populated by specific task sets, which include the concrete work to be done, the deliverables produced, and quality assurance points. This structure provides a foundation for organizing and executing the activities necessary for building high-quality software.

252. (A) Component-Based Development
Component-Based Development is a specialized software process model centered on the integration of pre-existing, reusable software components rather than building systems from scratch. This approach emphasizes the assembly of modular elements to construct complex applications efficiently. By focusing on established software blocks, development teams can improve productivity and reliability through the strategic reuse of verified components.

253. (C) To provide a template for solving a commonly occurring problem during development
A process pattern acts as a standardized template within a software framework to address recurring challenges encountered during development. It describes a proven solution for a specific problem within a given context, offering a structured set of activities or tasks. Implementing these patterns allows teams to utilize established best practices for solving common issues efficiently and consistently.

254. (D) Their use of mathematical notations for system specification and verification
Formal methods utilize rigorous mathematical notations to specify, develop, and verify software systems with high precision. This methodology provides a superior level of assurance by mathematically proving that a system meets its requirements. Unlike empirical approaches, these techniques focus on formal verification to ensure correctness, making them suitable for critical systems where failure-free operation is mandatory.

255. (C) They define a specific, predetermined workflow and set of process elements
Prescriptive process models provide a highly structured approach by defining a specific, predetermined workflow of activities, actions, and tasks. Unlike more flexible or general frameworks, these models establish a predictable sequence of process elements to be followed during software construction. This disciplined structure aims to bring order and clarity to the development lifecycle through established, sequential steps.

256. (B) Scrum
Scrum is an agile framework that manages complex software development by organizing work into fixed-length iterations called Sprints. Each Sprint represents a time-boxed period

Downloaded from www.pdf.tube

where specific tasks from the backlog are completed to deliver working software. This iterative approach facilitates continuous progress through structured events like planning and reviews, ensuring that the team remains focused on short-term goals.

257. (D) Tailoring of development processes to suit the specific context of the project and team
The Crystal family of methodologies is based on the principle that software development processes should be flexible and context-dependent. It suggests that different projects require unique policies and practices based on factors like team size and project criticality. By tailoring the development approach to the specific project environment, teams can optimize their efficiency and overall software quality.

258. (A) Employing techniques such as pair programming and continuous integration
Extreme Programming incorporates a set of specific technical and collaborative practices to improve software quality and responsiveness. Key elements of this agile methodology include pair programming, where two developers work together, and continuous integration, which ensures code is merged frequently. These practices emphasize technical excellence and sustainable development paces, helping teams maintain high standards throughout the software construction process.

259. (C) Incremental delivery of software and adaptability to change
Agile process models prioritize adaptability and the frequent delivery of functional software through incremental development cycles. This philosophy emphasizes responding to changing requirements and collaborating closely with customers rather than following rigid plans. By delivering software in small, working increments, teams can validate progress continuously and ensure the final product aligns with evolving stakeholder needs and business goals.

260. (A) Speculate, Collaborate, Learn
The Adaptive Software Development model relies on a cycle of speculating, collaborating, and learning to manage the uncertainty of complex projects. This iterative process allows development teams to adjust their approach based on new findings and stakeholder feedback. By emphasizing adaptability over rigid planning, the model ensures that the software evolves effectively to meet the changing needs of the business.

261. (D) A non-functional requirement, as it specifies a performance constraint
Non-functional requirements describe the quality attributes and operational constraints of a system rather than its specific behaviors. A performance-based constraint specifying processing speed illustrates how the software must perform its tasks. These criteria are essential for ensuring that the software meets efficiency, reliability, and other critical performance-related stakeholder expectations, providing a high-quality experience for users throughout its operation.

262. (B) To describe sequences of interactions between an actor and the system
Use cases are developed during the requirements phase to capture functional requirements from the perspective of the end-user. They describe the specific sequences of interactions between an actor, such as a user or external system, and the software to achieve a goal. This modeling technique helps developers understand system behavior and ensures that the requirements reflect actual user needs.

263. (C) The Software Requirement and Specification (SRS) document
The Software Requirement and Specification document serves as the official, detailed agreement between the development team and stakeholders. It records all functional and non-functional requirements, providing a foundation for the subsequent design and implementation phases. As the final output of the requirements process, this document acts as a definitive contract that outlines exactly what the software system will deliver.

264. (D) It is a quality assurance activity to check requirements for clarity and completeness
A requirements review is a formal quality assurance activity where stakeholders and developers meticulously examine the documented requirements. The primary goal is to identify and resolve inconsistencies, ambiguities, or omissions before the design phase begins. By scrutinizing

NODIA

requirements for clarity, completeness, and feasibility, the team ensures a solid foundation for building a high-quality software product that meets user needs.

265. (A) Functional requirements describe system behaviors, while non-functional requirements describe quality attributes
Functional requirements define the specific actions and behaviors the system must exhibit, such as data processing or report generation. Conversely, non-functional requirements specify quality attributes like performance, security, and usability. While functional requirements explain what the system does, non-functional requirements dictate how the system operates within its constraints, ensuring the software is both useful and effective for the user.

266. (B) Modularity
Modularity is a fundamental design principle that involves partitioning a complex software system into separate, independent components called modules. Each module is designed to address a specific part of the overall problem, allowing for better management of complexity. This division enables developers to work on different sections of the code in parallel while improving the overall maintainability and reusability.

267. (B) The elements within a module are functionally related and work towards a common purpose
High cohesion indicates that the elements contained within a single software module are closely related and focused on a single, well-defined task. In a good design, a cohesive module performs a specific function where all its internal parts work together toward a common goal. This attribute simplifies understanding and maintenance, as the module's purpose is clear and concentrated.

268. (A) High cohesion and low coupling are indicators of high functional independence
Functional independence is a qualitative measure of how well a module performs a specific task without relying heavily on other parts of the system. It is primarily assessed through cohesion and coupling. High cohesion ensures internal unity, while low coupling minimizes external dependencies. Achieving both characteristics leads to independent modules that are easier to develop, test, and modify without affecting others.

269. (C) Restricting access to the internal details of a module
Information hiding is a design principle that restricts access to the internal data structures and procedural logic of a module. By making these details inaccessible to other components that do not need them, the principle reduces complexity and prevents unintended dependencies. This approach ensures that changes made to a module's internal implementation do not negatively impact the rest of the system.

270. (D) High coupling
High coupling occurs when there is a significant degree of interdependence between different software modules. In such designs, modules are tightly connected, meaning that a modification in one component often requires extensive changes in several others. This ripple effect makes the system difficult to maintain and test, as the lack of independence increases the risk of introducing new defects.

271. (C) Architectural Design
Architectural design establishes the high-level framework of a software system by identifying the major components and their properties. It defines how these individual elements relate and interact with one another to form a cohesive structure. This activity serves as the blueprint for the entire system, providing a foundation for more detailed design tasks and ensuring that the overall requirements are met.

272. (B) Data Design
Data design is the process of translating data objects identified during the analysis phase into concrete data structures and database schemas. This activity focuses on defining how the software will organize, store, and manipulate information. By establishing these models early, designers ensure that the system can efficiently handle data attributes and relationships, which is crucial for the software's functional performance.

273. (D) Creating the detailed internal logic and algorithms for individual modules
Component-level design focuses on the internal implementation of individual software modules

Click Here to Buy book on Amazon

defined during the architectural phase. It involves creating the detailed procedural logic, algorithms, and data structures necessary for each component to perform its specific function. This detailed specification serves as the final blueprint for writing the actual source code, ensuring that every module operates correctly and meets its goals.

274. (A) Architectural Design -> Data Design / User Interface Design -> Component Level Design
Software design follows a logical progression starting from high-level structures to detailed internal logic. Architectural design first establishes the overall system framework. Subsequently, data design and user interface design address the information models and user interactions. Finally, component-level design provides the detailed procedural logic for individual modules. This sequence ensures that the software is built on a solid, well-defined structural foundation.

275. (C) User Interface Design
User interface design focuses on the interaction between the end-user and the software system. It encompasses the visual look and feel of the application, aiming to create an intuitive and aesthetically pleasing experience. By designing efficient interaction models and layouts, this activity ensures that users can navigate the software easily and perform their tasks effectively without unnecessary confusion or effort.

276. (A) Product revision
McCall's Quality Factors categorize software attributes into three perspectives to assess quality from different angles. Maintainability is classified under the product revision perspective, which focuses on the ease with which software can be modified. This includes the effort required to locate and fix errors or update the system, ensuring the software can evolve to meet changing requirements over its lifecycle.

277. (D) QA is process-oriented, focusing on defect prevention; QC is product-oriented, focusing on defect detection
Quality Assurance is a process-oriented discipline focused on preventing defects by establishing standards and auditing development procedures. In contrast, Quality Control is a product-oriented activity that involves detecting and removing faults in the actual software through testing and reviews. While QA aims to build quality into the process, QC ensures the final product meets specified standards by finding existing errors.

278. (B) Cost-effectiveness
The ISO 9126 standard identifies six primary characteristics for evaluating software quality: functionality, reliability, usability, efficiency, maintainability, and portability. While factors like cost-effectiveness are important for project management, they are not part of this specific technical quality framework. Each of the six standard characteristics is further divided into sub-characteristics to allow for a detailed and objective assessment of software quality.

279. (C) Quality Control (QC)
Quality control involves activities aimed at identifying and correcting defects in the final software products. Conducting a code review to find and fix errors is a classic example of a product-oriented quality control task. By scrutinizing the code for faults, the team ensures that the individual components satisfy the design specifications before the software proceeds to the next development stage.

280. (D) The probability of failure-free operation for a specified period in a specified environment
Software reliability is a quantitative measure defined as the probability that a system will operate without failure for a specific period within a given environment. This metric assesses the stability of the software under operational conditions. High reliability indicates that the software is dependable and performs its intended functions consistently, minimizing the likelihood of unexpected crashes or errors during actual usage.

281. (A) Risk Mitigation, Risk Monitoring, and Risk Management
The Risk Mitigation, Monitoring, and Management plan is a core component of proactive project management. It involves creating strategies to reduce the impact of risks, tracking the status of identified risks throughout the project, and executing management activities to control these threats. Together, these three elements provide a structured framework

for handling potential problems before they negatively affect the project.

282. (B) Risk analysis
Risk analysis is the systematic process of evaluating identified risks to determine their probability of occurrence and potential impact on the project. By estimating these factors, project managers can prioritize risks and allocate resources effectively to address the most significant threats. This analytical step is crucial for developing an informed mitigation strategy and ensuring that project goals remain achievable despite uncertainties.

283. (B) To track the status of identified risks and watch for triggers to invoke contingency plans
Risk monitoring is the continuous process of observing identified project risks to assess their current status and watch for specific triggers. These triggers indicate that a potential problem is becoming an actual issue, requiring the deployment of contingency plans. By maintaining ongoing surveillance of risks, the project manager can respond quickly and effectively, preventing minor issues from escalating into major failures.

284. (A) Risk mitigation
Risk mitigation involves developing proactive strategies to decrease the likelihood or the consequences of potential problems. One common mitigation approach is the creation of contingency plans that are ready to be implemented if a risk occurs. By planning these responses in advance, teams can reduce the overall impact on the project, ensuring that they are prepared for various unexpected scenarios.

285. (C) Because a functionally correct product may be unusable if its reliability is low
Managing software reliability as a risk is essential because a product that functions correctly but fails frequently is practically useless to stakeholders. If a system cannot maintain stable operation in its intended environment, its overall quality is compromised. Treating reliability as a potential risk ensures that developers focus on creating a robust product that remains operational and dependable under real-world conditions.

286. (D) LOC is a direct measure of code volume, while FP is an indirect measure based on functionality
Lines of Code and Function Points are two distinct metrics used for software sizing. Lines of Code is a direct measurement based on the total volume of program text. Conversely, Function Points provide an indirect measure based on the functionality delivered to the user, such as inputs and outputs. While LOC depends on implementation, FP focuses on the user's external perspective.

287. (C) To predict project effort and duration based on software size
The Constructive Cost Model is an algorithmic estimation tool used to predict the effort and duration required for a software project. It utilizes mathematical formulas based on historical data to provide estimates for costs and schedules. By applying these models, project managers can gain a realistic understanding of resource requirements, helping them plan the development process more accurately and effectively.

288. (B) The estimated size of the software in Kilo Lines of Code
The COCOMO model primarily uses the estimated size of the software, measured in Kilo Lines of Code, as its main input for calculating project effort. This size estimate is plugged into regression formulas derived from past projects to determine the man-months and total time needed for completion. Accurate sizing in KLOC is therefore critical for obtaining reliable effort and duration predictions.

289. (D) To visualize the project schedule, tasks, and their durations
A time-line chart, such as a Gantt chart, is a project management tool used to visualize the schedule and progression of tasks. It displays the start and end dates for various activities, their durations, and the dependencies between them. By using these charts, managers can track progress, manage personnel assignments, and ensure that the project remains on its planned timeline.

290. (A) Software Sizing -> Effort Estimation -> Project Scheduling
The logical sequence of project management activities begins with software sizing to quantify the scope of the work. This sizing data is then

Downloaded from www.pdf.tube

(D) It builds a solution incrementally and abandons it if it becomes invalid.

17. The N-Queens problem, where the goal is to place N chess queens on an N×N chessboard so that no two queens threaten each other, is a classic example of a problem solved using:
(A) Divide and Conquer
(B) Backtracking
(C) Dynamic Programming
(D) Greedy approach

18. Consider the following statements regarding graph traversal algorithms.
I. Breadth-First Search (BFS) is guaranteed to find the shortest path in an unweighted graph.
II. Depth-First Search (DFS) uses a stack data structure for its implementation.
III. BFS explores the graph layer by layer.
IV. The time complexity of both BFS and DFS on an adjacency list representation is O(V+E).
V. DFS is suitable for finding connected components in a graph.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) II, III, IV, and V only
(C) I, II, III, and IV only
(D) I, II, III, IV, and V

19. In a threaded binary tree, a NULL pointer in the original binary tree is replaced by a pointer to what?
(A) The root of the tree
(B) A special header node
(C) The in-order successor or in-order predecessor of the node
(D) The parent of the node

20. A Binary Search Tree (BST) has a specific ordering property. What is it?
(A) All nodes in the left subtree are greater than the root, and all nodes in the right subtree are smaller.
(B) The value of the root node is the average of its left and right children.
(C) The tree is always a complete binary tree.
(D) For any given node, all values in its left subtree are less than the node's value, and all values in its right subtree are greater.

21. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : An AVL tree is a self-balancing binary search tree.
Reason (R) : In an AVL tree, the heights of the two child subtrees of any node differ by at most one.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

22. What is the primary advantage of a B+ Tree over a B Tree for implementing the indexing of a database?
(A) B+ Trees are generally shorter in height than B Trees for a fixed number of keys.
(B) Data pointers are stored only in leaf nodes, making sequential access much faster.
(C) The process of deleting nodes is significantly simpler to implement in B+ Trees.
(D) B+ Trees utilize a higher branching factor, which helps in reducing memory usage.

23. Which data structure is specifically designed to efficiently support set operations like union, intersection, and difference?
(A) A stack
(B) A queue
(C) A simple array
(D) A hash table or a balanced binary search tree

24. Consider the following statements regarding the performance analysis of algorithms.
I. Big-O notation provides an asymptotic upper bound on the growth rate of a function.
II. Big-Omega (Ω) notation provides an asymptotic lower bound.
III. Big-Theta (Θ) notation provides an asymptotic tight bound.
IV. The space complexity of an algorithm measures the amount of time it takes to run.
Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II and III only
(C) III and IV only
(D) I, II, III and IV

299. (B) Change Control and Version Control
Software Configuration Management includes core functions such as change control and version control to manage the evolution of software products. Change control provides a mechanism for approving and tracking modifications, while version control maintains a history of different iterations of the software. Together, these functions ensure the integrity and traceability of the codebase, especially during complex development and reuse activities.

300. (A) Stress Testing
Stress testing is a non-functional testing technique used to evaluate the stability and robustness of a software system under extreme conditions. It involves subjecting the software to workloads that meet or exceed its specified limits to see how it handles resource exhaustion. This process helps identify the system's breaking point and ensures that it can recover gracefully from heavy operational loads.

Click Here to Buy book on Amazon

(C) Sorting
(D) Hashing

52. Consider a hash table of size 10 and a hash function h(key) = key mod 10. If linear probing is used to resolve collisions, where would the key 45 be placed if the keys 25, 35, and 15 are already in the table?
(A) Index 5
(B) Index 6
(C) Index 7
(D) Index 8

53. Consider the following characteristics of Dynamic Programming.
I. The problem can be divided into overlapping subproblems.
II. It uses a bottom-up approach, solving smaller subproblems first.
III. It is typically implemented using recursion with memoization or tabulation.
IV. It always leads to an exponential time complexity.
Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II, and III only
(C) I and III only
(D) All are correct

54. Which of the following is NOT a greedy algorithm?
(A) Prim's Algorithm
(B) 0/1 Knapsack Problem Algorithm
(C) Kruskal's Algorithm
(D) Dijkstra's Algorithm

55. The time complexity of finding an element in a balanced Binary Search Tree is:
(A) $O(1)$
(B) $O(n)$
(C) $O(n \log n)$
(D) $O(\log n)$

56. Match List I with List II.
List I: (Data Structure)
P. B Tree
Q. AVL Tree
R. Stack
S. Graph
List II: (Key Property/Use)
1. Height-balancing
2. Recursion implementation
3. Adjacency matrix/list
4. Multi-level indexing in files
Choose the correct option from those given below:
(A) P-1, Q-4, R-2, S-3
(B) P-4, Q-1, R-3, S-2
(C) P-4, Q-1, R-2, S-3
(D) P-2, Q-1, R-4, S-3

57. The process of visiting each node in a tree data structure exactly once is called:
(A) Searching
(B) Insertion
(C) Traversal
(D) Balancing

58. Consider the following statements.
Statement I : A B+ tree stores data records only in its leaf nodes.
Statement II : A B tree can store data records in both internal and leaf nodes.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

NODIA

59. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Backtracking algorithms are generally more efficient than brute-force enumeration.
Reason (R) : Backtracking prunes large parts of the solution space by abandoning paths that cannot lead to a solution.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

60. The worst-case number of comparisons required for binary search in a sorted array of 'n' elements is:
(A) n
(B) $\lfloor \log_2 n \rfloor + 1$
(C) $n/2$
(D) $n - 1$

Out (LIFO) storage manner.
(D) All elements must have a unique priority value within the structure.

9. Consider the following statements about tree data structures.
I. A tree is a non-linear data structure.
II. A collection of disjoint trees is called a forest.
III. Every non-empty tree has exactly one root node.
IV. A tree with N nodes always has N-1 edges.
Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II, and III only
(C) II, III, and IV only
(D) I, II, III, and IV

10. Identify the correct pair from the following, matching the algorithm with its average-case time complexity.
(A) Binary Search - $O(\log n)$
(B) Bubble Sort - $O(n \log n)$
(C) Merge Sort - $O(n^2)$
(D) Insertion Sort - $O(n \log n)$

11. What is the primary reason for using a good hashing function in a hash table?
(A) To guarantee that no two unique keys will map to the same internal array index.
(B) To ensure keys are stored in a sorted order for faster retrieval of elements.
(C) To increase the size of the hash table dynamically when it becomes nearly full.
(D) To minimize the number of collisions and ensure a uniform distribution of keys.

12. Arrange the following tree data structures in the chronological order of their conceptual development or popularization in computer science.
1. B Tree
2. Binary Tree
3. AVL Tree
4. B+ Tree
(A) 1, 2, 3, 4
(B) 2, 3, 1, 4
(C) 2, 1, 4, 3
(D) 3, 2, 1, 4

13. The Divide and Conquer strategy involves all of the following steps EXCEPT:
(A) Breaking the problem into smaller subproblems.
(B) Solving the subproblems recursively.
(C) Storing the results of subproblems to avoid re-computation.
(D) Combining the solutions to subproblems to get the final solution.

14. Consider the following statements.
Statement I : Dynamic Programming is an algorithmic technique that solves complex problems by breaking them down into simpler overlapping subproblems.
Statement II : The Fibonacci sequence calculation is a classic example where Dynamic Programming can be applied to optimize performance.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

NODIA

15. Match List I with List II.
List I: (Algorithm Design Technique)
P. Divide and Conquer
Q. Dynamic Programming
R. Greedy Algorithm
S. Backtracking
List II: (Problem Example)
1. 0/1 Knapsack Problem
2. N-Queens Problem
3. Merge Sort
4. Fractional Knapsack Problem
Choose the correct option from those given below:
(A) P-1, Q-3, R-2, S-4
(B) P-3, Q-1, R-4, S-2
(C) P-3, Q-4, R-1, S-2
(D) P-4, Q-1, R-2, S-3

16. Which of the following best describes a Greedy Algorithm?
(A) It explores all possible solutions and chooses the best one for the problem.
(B) It makes a locally optimal choice at each stage to find a global optimum.
(C) It divides the problem into subproblems, solves them, and combines results.

used in effort estimation to determine the required resources, time, and budget. Finally, project scheduling distributes that estimated effort across a timeline. Following this order ensures that the schedule is based on realistic estimates derived from project size.

291. (C) Does the software conform to its design specification?
Verification is the process that ensures software products satisfy conditions imposed at the beginning of a specific development phase. It focuses on checking if the system conforms to its established design specifications and technical requirements. By confirming that development proceeds according to plan, teams ensure that the final product is built correctly and adheres to all technical standards.

292. (A) An Error leads to a Fault, which may cause a Failure
The relationship between errors, faults, and failures explains how defects manifest in software. A human mistake during development creates a fault or bug within the source code. When this fault is triggered during execution, it results in a failure, which is an observable deviation from the system's expected behavior. Understanding this process helps developers identify, track, and resolve underlying issues.

293. (D) Unit Testing
Unit testing is the process of verifying the smallest individual components or modules of a software system in isolation. The primary objective is to ensure that each unit functions correctly and meets its specific requirements before being integrated with other parts. By testing modules independently, developers can easily identify and fix defects within the procedural logic early in the development lifecycle.

294. (B) Alpha testing is done at the developer's site, while beta testing is done at the end-user's site
Alpha and beta testing are distinct stages of user acceptance testing used to validate software. Alpha testing is conducted by internal teams at the developer's site to find initial defects. In contrast, beta testing is performed by actual users in their own real-world environments. This allows the software to be evaluated under diverse conditions before its final release to the general public.

295. (C) Expose defects in the interfaces and interactions between components
Integration testing is performed to identify defects that occur when different software components interact with one another. While individual modules may work correctly in isolation, issues often arise in their interfaces or data exchanges when combined. The goal of this testing level is to ensure that the integrated system components function together smoothly and fulfill the requirements of the overall architecture.

296. (D) White-box testing
White-box testing is a technique where test cases are derived from the internal logic and structure of the software code. One specific method involves using the program's control flow graph to identify and execute every independent path. By examining the code directly, testers can ensure that all conditions and internal paths are thoroughly exercised, improving the overall reliability of the software system.

297. (A) To ensure that modifications have not introduced new defects in existing features
Regression testing is conducted after software modifications, such as bug fixes or feature updates, to ensure that existing functionality remains intact. The primary purpose is to verify that recent changes have not inadvertently introduced new defects into previously working sections of the code. This continuous verification process is essential for maintaining software integrity and quality during the evolution of the system.

298. (B) Black-box testing focuses on functionality, while White-box testing focuses on internal structure
Black-box testing and white-box testing differ in their focus and information requirements. Black-box testing evaluates the software based on functional requirements without knowledge of the internal code. In contrast, white-box testing focuses on the internal structure and logic, requiring a detailed understanding of the source code. While one checks what the system does, the other verifies how the code is structured.

Downloaded from www.pdf.tube

25. The recurrence relation that arises from the Merge Sort algorithm is:
(A) $T(n) = T(n-1) + O(n)$
(B) $T(n) = T(n/2) + O(1)$
(C) $T(n) = 2T(n/2) + O(n)$
(D) $T(n) = 2T(n-1) + O(1)$

26. Which of the following is NOT a characteristic of the Branch and Bound algorithm design technique?
(A) It does not use any bounding function to prune the search space.
(B) It is used for optimization problems.
(C) It explores the state space tree using breadth-first search or best-first search.
(D) It maintains the best solution found so far to prune subtrees.

27. Match List I with List II.
List I: (Complexity Class/Concept)
P. P Class
Q. NP Class
R. NP-complete
S. Reducibility
List II: (Description)
1. The class of problems for which a solution can be verified in polynomial time.
2. A problem in NP to which all other NP problems can be reduced in polynomial time.
3. The class of problems that can be solved in polynomial time by a deterministic Turing machine.
4. A technique to transform one problem into another.
Choose the correct option from those given below:
(A) P-3, Q-1, R-2, S-4
(B) P-1, Q-3, R-2, S-4
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-2, S-3

28. What does the Lower Bound Theory in algorithm analysis primarily concern itself with?
(A) Finding the best-case running time for a specific algorithm.
(B) Establishing the minimum amount of work any algorithm must perform to solve a given problem.
(C) Calculating the average-case complexity of sorting algorithms.
(D) Determining the upper limit on the resources required by an algorithm.

29. A comparison tree for sorting 'n' elements has how many leaves at a minimum?
(A) n
(B) n^2
(C) $n!$
(D) 2^n

30. Consider the following statements about advanced algorithm concepts.
I. Randomized algorithms use a degree of randomness as part of their logic.
II. Approximation algorithms are used for NP-hard problems to find a near-optimal solution in polynomial time.
III. Parallel algorithms are designed to run on machines with multiple processors.
IV. The Fast Fourier Transform (FFT) is an algorithm to compute the Discrete Fourier Transform (DFT).
V. All randomized algorithms are guaranteed to produce the correct result.
Which of the above statement(s) is/are correct?
(A) I, II, and V only
(B) I, II, III, and IV only
(C) II, III, and V only
(D) All are correct

31. If a problem A can be reduced to problem B in polynomial time ($A \leq_p B$), and B is in P class, what can be concluded about problem A?
(A) A is also in P class.
(B) A is in NP class.
(C) A is NP-hard.
(D) No conclusion can be drawn about A.

32. The Knuth-Morris-Pratt (KMP) algorithm is a well-known example of which of the following?
(A) A polynomial arithmetic algorithm
(B) A string matching algorithm
(C) A number theoretic algorithm
(D) A graph traversal algorithm

33. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : In a B-Tree of order m, each internal node can have at most m children.
Reason (R) : Each internal node in a B-Tree can

NODIA

CHAPTER 07

Data Structures and Algorithms

1. Which data structure operates on a Last-In, First-Out (LIFO) principle?
(A) Queue
(B) Stack
(C) Linked List
(D) Tree

2. Consider the following statements.
Statement I : Arrays provide a random access capability to elements using an index.
Statement II : The size of an array is dynamic and can be changed at runtime without reallocating memory.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

3. Match List I with List II.
List I: (Data Structure)
P. Stack
Q. Queue
R. Priority Queue
S. Sparse Matrix
List II: (Primary Application)
1. Job scheduling in an operating system
2. Representing a grid with mostly zero values
3. Function call management
4. First-come, first-served resource allocation
Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-2, S-1
(C) P-3, Q-1, R-4, S-2
(D) P-1, Q-2, R-3, S-4

4. A linear data structure where insertions happen at one end, called REAR, and deletions happen at the other end, called FRONT, is known as a:
(A) Stack
(B) Deque
(C) Priority Queue
(D) Queue

5. Which of the following statements about a singly linked list is INCORRECT?
(A) Each node contains a data field and a pointer to the next node.
(B) It allows for efficient insertion and deletion of elements.
(C) Any node can be accessed directly using an index in O(1) time.
(D) The last node in the list points to NULL.

6. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The worst-case time complexity of Quick Sort is $O(n^2)$.
Reason (R) : The worst-case occurs when the pivot element is always the smallest or largest element in the partition.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

7. A company stores a large 1000x1000 matrix representing a social network, where most people are not connected. To save memory, which data structure is most appropriate to store this matrix?
(A) A standard 2D array representation.
(B) A threaded binary tree structure.
(C) A hash table with chaining method.
(D) A sparse matrix representation.

8. What is the defining characteristic of a Priority Queue?
(A) Elements are processed based on a priority value, not insertion order.
(B) Elements are processed in the order they are inserted into the queue.
(C) Elements are processed in a Last-In, First-

Click Here to Buy book on Amazon

List II: (Problem Solved)
1. All-pairs shortest paths
2. Single-source shortest paths (unweighted graph)
3. Minimum Spanning Tree
4. Single-source shortest paths (non-negative weights)

Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-3, Q-2, R-1, S-4
(C) P-4, Q-3, R-1, S-2
(D) P-1, Q-4, R-3, S-2

43. A Monte Carlo algorithm is a type of randomized algorithm that:
(A) always produces a correct answer, but its running time varies randomly.
(B) uses random sampling to approximate a definite integral.
(C) combines solutions from multiple parallel computations.
(D) has a deterministic running time but has a small probability of producing an incorrect answer.

44. Which of the following is NOT an application of arrays?
(A) Implementing other data structures like stacks and queues.
(B) Managing function calls in a recursive program.
(C) Storing a list of employee records.
(D) Representing a lookup table.

45. Consider the following statements about linked lists.
I. A doubly linked list allows traversal in both forward and backward directions.
II. A circular linked list has its last node pointing to the first node instead of NULL.
III. Insertion at the beginning of a singly linked list takes O(n) time.
IV. Deleting a node from a doubly linked list is more efficient than from a singly linked list if a pointer to the node to be deleted is given.
V. A linked list requires contiguous memory allocation for its nodes.

Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, and IV only
(C) III, IV, and V only
(D) I, II, IV, and V only

46. What is the postfix expression for the infix expression `A + B * (C - D) / E`?
(A) A B + C D - * E /
(B) A B C D * - E / +
(C) A B + C D - E / *
(D) A B C D - * E / +

47. The height of a perfectly balanced binary tree with N nodes is approximately:
(A) $\log_2 N$
(B) N
(C) $N/2$
(D) $N \log_2 N$

48. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The problem of determining if a number is prime is in the NP class.
Reason (R) : If a number is given along with one of its factors, we can verify in polynomial time that it is composite.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

NODIA

49. Which sorting algorithm is considered stable?
(A) Quick Sort (standard implementation)
(B) Merge Sort
(C) Heap Sort
(D) Selection Sort

50. The recurrence relation $T(n) = T(n-1) + O(1)$ describes the time complexity of which operation?
(A) Traversing a linked list to find the last element.
(B) Binary search in a sorted array.
(C) Building a heap from an array.
(D) Merge sort on an array of size n.

51. The process of arranging records in a specific order is known as:
(A) Searching
(B) Merging

contain at most m-1 keys.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

34. Which of the following is an example of a problem that is known to be in P class?
(A) Finding the Minimum Spanning Tree of a graph
(B) The Traveling Salesperson Problem (TSP)
(C) The 3-SAT problem
(D) The vertex cover problem

35. The Fast Fourier Transform (FFT) algorithm improves the complexity of polynomial multiplication from $O(n^2)$ to:
(A) $O(n)$
(B) $O(\log n)$
(C) $O(n \log n)$
(D) $O(n^{1.5})$

36. Arrange the following sorting algorithms based on their worst-case time complexity in ascending order (from best to worst).
1. Merge Sort
2. Bubble Sort
3. Quick Sort
4. Heap Sort

(A) 1, 4, 2, 3
(B) 4, 1, 3, 2
(C) 1, 3, 4, 2
(D) 1, 4, 3, 2

37. Which of the following statements about graph algorithms is INCORRECT?
(A) Dijkstra's algorithm always finds the shortest path in graphs with non-negative edge weights.
(B) Kruskal's algorithm for MST follows a greedy approach.
(C) Breadth-First Search (BFS) uses a stack for its implementation.
(D) Depth-First Search (DFS) can be used to detect cycles in a graph.

38. The problem of finding the maximum flow from a source 's' to a sink 't' in a network is a classic graph problem. Which algorithm is commonly used to solve this?
(A) Dijkstra's algorithm
(B) Ford-Fulkerson algorithm
(C) Prim's algorithm
(D) Bellman-Ford algorithm

39. Consider the following statements about B* Trees.
I. A B* Tree is a variant of a B-Tree.
II. B* Trees ensure that each internal node is at least 2/3 full.
III. Splitting of a node is delayed by redistributing keys to a sibling node.
IV. B* Trees are less space-efficient than standard B-Trees.

Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II, and III only
(C) II and IV only
(D) I, III, and IV only

40. Consider the following statements.
Statement I : An algorithm with a time complexity of $O(n)$ is asymptotically faster than an algorithm with a time complexity of $O(\log n)$.
Statement II : The function $f(n) = n \log n$ has a higher growth rate than $f(n) = n$.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

41. In the context of approximation algorithms, what does an approximation ratio of ρ signify?
(A) The algorithm runs ρ times faster than the known optimal algorithm.
(B) The probability of finding the exact optimal solution is given by ρ.
(C) The solution found is guaranteed to be within a factor of ρ of the optimum.
(D) The algorithm uses ρ percent less memory than a standard exact algorithm.

42. Match List I with List II.
List I: (Algorithm)
P. Prim's Algorithm
Q. Dijkstra's Algorithm
R. Floyd-Warshall Algorithm
S. Breadth-First Search

Downloaded from www.pdf.tube

61. Which notation provides an asymptotic tight bound for the growth rate of a function in algorithm analysis?
(A) Big O
(B) Big Omega (Ω)
(C) Big Theta (Θ)
(D) Little o

62. Arrange the following algorithm design paradigms in the general order of their formalization and widespread application.
1. Greedy Algorithms
2. Dynamic Programming
3. Divide and Conquer
4. Backtracking

(A) 3, 2, 1, 4
(B) 1, 3, 2, 4
(C) 4, 3, 2, 1
(D) 3, 1, 4, 2

63. The maximum number of nodes in a binary tree of height 'h' (where root is at height 0) is:
(A) $2^{h+1} - 1$
(B) $2^h - 1$
(C) 2^h
(D) $2h + 1$

64. For an undirected graph with 'n' vertices and 'e' edges, the sum of degrees of all vertices is equal to:
(A) e
(B) 2e
(C) n-1
(D) e+1

65. Identify the correct statements regarding Complexity Theory.
I. A problem is in class P if it can be solved in polynomial time.
II. A problem is in class NP if a given solution can be verified in polynomial time.
III. Every problem in P is also in NP.
IV. The relationship P = NP is a proven theorem.
V. An NP-complete problem is both in NP and is NP-hard.

Which of the above statement(s) is/are correct?
(A) I, II, III, and V only
(B) I, II, and IV only
(C) I, III, IV, and V only
(D) II, III, and V only

NODIA

66. What is the best-case time complexity of the standard Bubble Sort algorithm?
(A) $O(n)$
(B) $O(1)$
(C) $O(n \log n)$
(D) $O(n^2)$

67. A graph in which there is a path between every pair of distinct vertices is called:
(A) A complete graph
(B) A cyclic graph
(C) A connected graph
(D) A directed acyclic graph

68. What is the balance factor of a node in an AVL tree?
(A) The difference between the number of nodes in the left and right subtrees.
(B) The maximum height of the node's subtrees.
(C) The number of children of the node.
(D) The difference between the height of the left subtree and the height of the right subtree.

69. In polynomial arithmetic, representing polynomials as arrays of coefficients allows for efficient implementation of:
(A) Addition and subtraction in $O(n)$ time.
(B) Multiplication in $O(n \log n)$ time using FFT.
(C) All of the above.
(D) Evaluation at a point using Horner's method.

70. Which of the following statements about Minimum Spanning Trees (MST) is INCORRECT?
(A) An MST of a graph with V vertices contains exactly V-1 edges.
(B) The MST is guaranteed to contain the shortest path between any two vertices in the graph.
(C) There can be multiple MSTs for a given graph if edge weights are not unique.
(D) Both Prim's and Kruskal's algorithms can find an MST.

71. Match List I with List II.
List I: (Sorting Method)
P. Selection Sort
Q. Insertion Sort
R. Quick Sort
S. Heap Sort

List II: (Basic Idea)
1. Build a max-heap and repeatedly extract the maximum element.
2. Partition the array and recursively sort the partitions.
3. Repeatedly find the minimum element from the unsorted part and put it at the beginning.
4. Build the final sorted array one item at a time.

Choose the correct option from those given below:
(A) P-3, Q-4, R-2, S-1
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-1, R-2, S-4
(D) P-2, Q-4, R-3, S-1

72. The concept of using randomness to solve problems that might be deterministic but very slow is central to:
(A) Randomized Algorithms
(B) Approximation Algorithms
(C) Dynamic Programming
(D) Backtracking

73. Consider the following statements.
Statement I : Reducibility is a key concept used to prove that a problem is NP-hard.
Statement II : To prove problem X is NP-hard, we can show that a known NP-complete problem can be reduced to X in polynomial time.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

74. A full binary tree with 'L' leaves has a total of how many nodes?
(A) 2L - 1
(B) L + 1
(C) 2L
(D) $\log_2 L$

75. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Dijkstra's algorithm may not produce a correct shortest path tree if the graph has negative edge weights.
Reason (R) : The greedy choice of Dijkstra's algorithm (picking the unvisited vertex with the smallest distance) can be invalidated by a negative edge encountered later.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

76. The problem of multiplying two large n-bit integers can be solved more efficiently than the classical $O(n^2)$ method using a Divide and Conquer algorithm like:
(A) Strassen's matrix multiplication
(B) Karatsuba algorithm
(C) Fast Fourier Transform
(D) Euclid's algorithm

77. Which data structure is most suitable for implementing a Breadth-First Search (BFS) in a graph?
(A) Stack
(B) Priority Queue
(C) Queue
(D) Array

78. Which of the following statements about Hashing are true?
I. The goal of a hash function is to distribute keys uniformly across the hash table.
II. A collision occurs when two different keys map to the same hash value.
III. Open addressing and chaining are two common methods for collision resolution.
IV. A perfect hash function is one that causes no collisions for a given set of keys.

Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II, and III only
(C) II, III, and IV only
(D) I, II, III, and IV

79. The time complexity to delete a node from a doubly linked list, given a pointer to the node, is:
(A) $O(1)$
(B) $O(\log n)$
(C) $O(n)$
(D) $O(n \log n)$

80. What is the lower bound for comparison-based sorting algorithms in the worst case?

Click Here to Buy book on Amazon

(A) $\Omega(n)$
(B) $\Omega(n \log n)$
(C) $\Omega(\log n)$
(D) $\Omega(n^2)$

81. Consider a queue implemented using an array of size MAX_SIZE. The condition `(rear + 1) % MAX_SIZE == front` signifies what state?
(A) The queue is empty.
(B) The queue is full.
(C) The queue has exactly one element.
(D) The queue is half-full.

82. In the context of Branch and Bound, what is a bounding function?
(A) A function that estimates a lower bound (for minimization) or an upper bound (for maximization) on the solution cost within a subtree.
(B) A function that calculates the exact cost of the optimal solution in a subtree.
(C) A function that determines the maximum depth of the state space tree.
(D) A function that counts the number of nodes in the state space tree.

83. Arrange the following search algorithms in the order of their development/popularization.
1. Binary Search
2. Linear Search
3. Fibonacci Search
4. Interpolation Search

(A) 2, 1, 4, 3
(B) 1, 2, 3, 4
(C) 2, 1, 3, 4
(D) 2, 3, 1, 4

84. The solution to the recurrence relation $T(n) = T(n/2) + O(1)$ is:
(A) $O(n)$
(B) $O(n \log n)$
(C) $O(n^2)$
(D) $O(\log n)$

85. Which of the following is an example of a Number Theoretic Algorithm?
(A) Kruskal's algorithm for minimum spanning trees.
(B) The Rabin-Karp string matching algorithm.
(C) The Sieve of Eratosthenes for finding prime numbers.
(D) The Fast Fourier Transform.

86. Identify the correct statements regarding graph representations.
I. An adjacency matrix for a graph with V vertices requires $O(V^2)$ space.
II. An adjacency list for a graph with V vertices and E edges requires $O(V + E)$ space.
III. Checking if an edge exists between two vertices is faster in an adjacency matrix than in an adjacency list.
IV. For a sparse graph, an adjacency list is more space-efficient than an adjacency matrix.
V. An adjacency matrix can only represent unweighted graphs.

Which of the above statement(s) is/are correct?
(A) I, II, III, and IV only
(B) I, II, and V only
(C) II, III, IV, and V only
(D) All are correct

87. Which of the following is the most appropriate data structure for simulating a waiting line, like customers at a ticket counter?
(A) Stack
(B) Tree
(C) Graph
(D) Queue

88. The process of finding the location of a specific element in a list of elements is called:
(A) Searching
(B) Sorting
(C) Merging
(D) Traversing

89. Match List I with List II.
List I: (Problem)
P. Finding a cycle in a graph
Q. Finding the shortest path in an unweighted graph
R. Solving the Tower of Hanoi
S. Evaluating a postfix expression

List II: (Data Structure/Algorithm)
1. Recursion
2. Stack
3. Depth-First Search (DFS)
4. Breadth-First Search (BFS)

Choose the correct option from those given below:
(A) P-3, Q-4, R-2, S-1

Downloaded from www.pdf.tube

(B) Bellman-Ford algorithm
(C) Breadth-First Search
(D) Prim's algorithm

109. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Dynamic programming solutions often have a polynomial time complexity.
Reason (R) : By storing and reusing the solutions to overlapping subproblems, dynamic programming avoids recomputing the same work repeatedly, which would otherwise lead to exponential complexity.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

110. Which statement about a sparse matrix is correct?
(A) It is a matrix where the number of non-zero elements is very high.
(B) It is a matrix that is best stored using a standard 2D array to save space.
(C) It is another name for an identity matrix.
(D) It is a matrix where most of the elements are zero.

111. The in-order traversal of a binary search tree results in visiting the nodes in:
(A) Ascending sorted order.
(B) Descending sorted order.
(C) The order of insertion.
(D) Level-by-level order.

112. Arrange the following steps of the Backtracking paradigm for solving a problem.
1. Check if the current partial solution can be extended to a complete solution.
2. If the path is promising, recursively explore it.
3. If the path is not promising, prune the branch (backtrack).
4. Define a state space that includes all possible solutions.

(A) 4, 1, 2, 3
(B) 1, 2, 3, 4
(C) 4, 2, 1, 3
(D) 1, 4, 2, 3

113. What is the worst-case time complexity for inserting an element into a Binary Search Tree (BST)?
(A) $O(1)$
(B) $O(n)$
(C) $O(\log n)$
(D) $O(n \log n)$

114. The problem of finding if a Boolean formula is satisfiable (SAT) is a canonical example of a problem that is:
(A) NP-complete
(B) In P
(C) Solvable in logarithmic time
(D) Not in NP

115. Consider the following properties of queues.
I. It is a FIFO (First-In, First-Out) data structure.
II. Enqueue is the operation to insert an element.
III. Dequeue is the operation to remove an element.
IV. A common application is simulating a waiting line.
Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II and III only
(C) II, III and IV only
(D) I, II, III and IV

116. In parallel computing, what is a common strategy for parallelizing searching algorithms on a sorted list?
(A) Each processor searches the entire list.
(B) The list is converted into a heap first.
(C) The search key is split among processors.
(D) The list is partitioned, and each processor searches one partition.

117. An approximation algorithm for the traveling salesman problem (TSP) provides a tour whose length is guaranteed to be no more than 1.5 times the optimal tour length. What is its approximation ratio?
(A) 1.0
(B) 2.0
(C) 1.5

NODIA

(C) $O(\log m + \log n)$
(D) $O(1)$

99. In the context of graph theory, what is a forest?
(A) A graph with no cycles.
(B) A tree where every node has at most one child.
(C) A collection of disjoint trees.
(D) Any subgraph of a tree.

100. The post-order traversal of a binary search tree gives the elements in:
(A) Ascending order
(B) Descending order
(C) The order they were inserted
(D) No particular order

101. Which of the following is NOT a type of B-Tree?
(A) B Tree
(B) B+ Tree
(C) B# Tree
(D) B* Tree

102. An algorithm solves a problem of size 'n' by recursively solving two subproblems of size n/2 and then combining the results in linear time. The recurrence relation for this is:
(A) $T(n) = T(n/2) + O(n)$
(B) $T(n) = 2T(n/2) + O(n)$
(C) $T(n) = 2T(n/2) + O(1)$
(D) $T(n) = T(n-1) + O(n)$

103. Which of the following statements about AVL trees are correct?
I. It is a self-balancing binary search tree.
II. The balance factor for any node is restricted to {-1, 0, 1}.
III. Rotations (like LL, RR, LR, RL) are performed to maintain balance after insertion or deletion.
IV. The height of an AVL tree with n nodes is guaranteed to be $O(\log n)$.
V. Searching in an AVL tree has a worst-case complexity of $O(n)$.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, IV, and V only
(D) All are correct

104. What is the primary purpose of a threaded binary tree?
(A) To enable traversal without using recursion or a stack.
(B) To make the tree balanced.
(C) To store more data in each node.
(D) To speed up insertion operations.

105. Consider the following statements.
Statement I : The 0/1 Knapsack problem can be solved optimally using a greedy approach.
Statement II : The Fractional Knapsack problem can be solved optimally using a greedy approach.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

106. Match List I with List II.
List I: (Asymptotic Notation)
P. $O(1)$
Q. $O(n)$
R. $O(n^2)$
S. $O(\log n)$
List II: (Common Name)
1. Logarithmic Time
2. Linear Time
3. Constant Time
4. Quadratic Time
Choose the correct option from those given below:
(A) P-3, Q-2, R-4, S-1
(B) P-3, Q-1, R-4, S-2
(C) P-1, Q-2, R-4, S-3
(D) P-2, Q-3, R-1, S-4

107. The process of converting an infix expression to a postfix expression is typically done using which data structure?
(A) Queue
(B) Linked List
(C) Tree
(D) Stack

108. Which graph algorithm is suitable for finding the shortest path in a weighted graph that contains negative-weight edges but no negative-weight cycles?
(A) Dijkstra's algorithm

Click Here to Buy book on Amazon

147. The concept of using memoization is a key feature of which algorithm design technique?
(A) Greedy approach
(B) Divide and Conquer
(C) Dynamic Programming (top-down approach)
(D) Backtracking

148. Which of the following statements correctly describe properties of sorting algorithms?
I. Bubble sort has a worst-case time complexity of $O(n^2)$.
II. Merge sort is a stable sort.
III. Quick sort's worst-case performance is worse than its average-case performance.
IV. Heap sort is an in-place sorting algorithm.
V. Selection sort performs the minimum number of swaps.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, IV, and V only
(C) I, II, III, IV, and V
(D) II, III, and V only

149. The Euclid's algorithm is a classic example of a:
(A) String Matching Algorithm
(B) Graph Algorithm
(C) Number Theoretic Algorithm
(D) Sorting Algorithm

150. A graph with no cycles is called an:
(A) Complete graph
(B) Acyclic graph
(C) Bipartite graph
(D) Regular graph

151. What is the primary purpose of Asymptotic Notation in algorithm analysis?
(A) To determine the exact running time of an algorithm.
(B) To count the number of statements in a program.
(C) To compare algorithms based on their implementation details.
(D) To describe the limiting behavior of a function when the argument tends towards a particular value or infinity.

152. In a B-Tree of minimum degree 't', what is the maximum number of keys an internal node can have?
(A) 2t - 1
(B) t - 1
(C) t
(D) 2t

153. Consider the following characteristics of a stack.
I. It is a LIFO (Last-In, First-Out) data structure.
II. PUSH is the operation to insert an element.
III. POP is the operation to remove an element.
IV. A common application is reversing a string.
Which of the above statement(s) is/are correct?
(A) I and II only
(B) I, II, and III only
(C) II, III, and IV only
(D) I, II, III, and IV

154. If an algorithm has a time complexity of $O(n^2)$, and the input size is doubled, the running time will increase by approximately a factor of:
(A) 4
(B) 2
(C) 8
(D) $\log 2$

155. Consider the following statements.
Statement I : Parallel algorithms can achieve a significant speedup by dividing a task among multiple processors.
Statement II : The speedup of a parallel algorithm is always equal to the number of processors used.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

156. Match List I with List II.
List I: (Tree Type)
P. Binary Search Tree
Q. AVL Tree
R. B+ Tree
S. Threaded Binary Tree
List II: (Defining Feature)
1. Null pointers replaced by threads to successor/predecessor
2. Used for database indexing, data in leaves
3. Keys in left subtree are smaller, keys in right are larger
4. Height-balanced binary search tree

NODIA

(D) 0.5

118. Which of the following areas does the Fast Fourier Transform (FFT) algorithm primarily apply to?
(A) String Matching Algorithms
(B) Polynomial Arithmetic
(C) Number Theoretic Algorithms
(D) Graph Algorithms

119. A data structure specifically designed to represent disjoint sets and perform find and union operations is called:
(A) An AVL Tree.
(B) A Priority Queue.
(C) A Sparse Matrix.
(D) A Union-Find or Disjoint Set Union (DSU) data structure.

120. The time and space complexities of an algorithm are evaluated based on:
(A) The size of the input.
(B) The speed of the computer it runs on.
(C) The number of lines of code.
(D) The programming language used.

121. Match List I with List II.
List I: (Algorithm)
P. Backtracking
Q. Dynamic Programming
R. Divide and Conquer
S. Greedy Method
List II: (Characteristic Feature)
1. Optimal substructure and overlapping subproblems
2. Makes locally optimal choice
3. State-space search with pruning
4. Breaks problem into independent subproblems
Choose the correct option from those given below:
(A) P-3, Q-1, R-4, S-2
(B) P-1, Q-3, R-4, S-2
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-2, S-3

122. Which of the following statements about space complexity is correct?
(A) It is the amount of memory space required by an algorithm, as a function of the input size.
(B) It is the amount of hard disk space required by the program.
(C) It is always less than the time complexity of the algorithm.
(D) It is a measure of the CPU time taken by the algorithm.

123. Consider the following statements.
Statement I : The Master Theorem can be used to solve recurrence relations of the form $T(n) = aT(n/b) + f(n)$.
Statement II : The Master Theorem can be applied to any recurrence relation without exception.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

124. The minimum number of nodes in an AVL tree of height 'h' is a function related to:
(A) The factorial function
(B) The Fibonacci numbers
(C) The powers of 2
(D) The Catalan numbers

125. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : For a very sparse graph, an adjacency list representation is generally preferred over an adjacency matrix.
Reason (R) : The space complexity of an adjacency list is O(V+E), while for an adjacency matrix it is O(V^2), which is inefficient when the number of edges E is much smaller than V^2.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

126. If a problem is NP-complete, it means that:
(A) There is no algorithm to solve it.
(B) It can be solved in polynomial time.
(C) It is not verifiable in polynomial time.
(D) It is one of the hardest problems in the NP class.

Downloaded from www.pdf.tube

127. Which sorting algorithm has the best performance for an array that is already sorted?
(A) Insertion Sort
(B) Quick Sort
(C) Selection Sort
(D) Heap Sort

128. Consider the various tree structures.
I. A binary tree node can have at most two children.
II. In a Binary Search Tree, the left child is always smaller and the right child is always larger than the parent.
III. An AVL tree is a balanced version of a Binary Search Tree.
IV. A B-Tree is a generalization of a Binary Search Tree that can have more than two children.
V. A Threaded Binary Tree facilitates traversals without recursion.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, IV, and V only
(D) I, II, III, IV, and V

129. The lower bound theory proves that any comparison-based sorting algorithm requires at least $\Omega(n \log n)$ comparisons. This is derived using:
(A) Recurrence relations
(B) Decision trees (Comparison trees)
(C) Amortized analysis
(D) Strassen's method

130. The operation of processing each element in a list is known as:
(A) Traversal
(B) Sorting
(C) Merging
(D) Insertion

131. Which of the following is NOT a characteristic of a B+ Tree?
(A) All leaves are at the same level.
(B) Data pointers are stored only in leaf nodes.
(C) Internal nodes store both keys and data records.
(D) Leaf nodes are linked together to provide sequential access.

132. The Bellman-Ford algorithm for single-source shortest paths has a time complexity of:
(A) O(E log V)
(B) O(V + E)
(C) O(V^2)
(D) O(V * E)

133. Identify the correct statements about P and NP classes.
I. P is the class of decision problems solvable in polynomial time.
II. NP is the class of decision problems verifiable in polynomial time.
III. It is proven that P is a proper subset of NP.
IV. If P=NP, then many hard problems like TSP and SAT would have efficient solutions.
Which of the above statement(s) is/are correct?
(A) I, II, and IV only
(B) I and II only
(C) I, III, and IV only
(D) All are correct

134. Which algorithm design technique builds a solution step-by-step, and discards a candidate solution as soon as it determines it cannot be completed to a valid solution?
(A) Dynamic Programming
(B) Backtracking
(C) Divide and Conquer
(D) Greedy method

135. The result of a Depth-First Search (DFS) traversal on a graph is often represented as a:
(A) Queue of visited nodes
(B) Shortest path tree
(C) Minimum spanning tree
(D) DFS Forest/Tree

136. Match List I with List II.
List I: (Task)
P. Finding connected components
Q. Finding shortest path in unweighted graph
R. Topological sorting
S. Finding minimum spanning tree
List II: (Suitable Algorithm)
1. BFS
2. Kruskal's or Prim's Algorithm
3. DFS
4. Kahn's Algorithm or DFS
Choose the correct option from those given below:

(A) P-3, Q-1, R-4, S-2
(B) P-1, Q-3, R-2, S-4
(C) P-3, Q-1, R-2, S-4
(D) P-4, Q-1, R-3, S-2

137. A circular queue is 'empty' under which condition (using variables 'front' and 'rear' as indices, with -1 for initial empty state)?
(A) `front == rear`
(B) `(rear + 1) % MAX_SIZE == front`
(C) `front == -1`
(D) `rear == MAX_SIZE - 1`

138. Consider the following statements.
Statement I : The height of a B-tree is kept low by allowing nodes to have many children.
Statement II : The height of a B-tree grows logarithmically with the number of keys.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

139. If you need to implement a undo feature in a text editor, which data structure would be most suitable?
(A) Queue
(B) Tree
(C) Graph
(D) Stack

140. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A greedy algorithm does not always guarantee a globally optimal solution.
Reason (R) : Greedy algorithms make choices that are locally optimal at each step, which may not lead to the overall best solution for all types of problems.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

141. The worst-case time complexity of building a heap (heapify) of 'n' elements is:
(A) $O(\log n)$
(B) $O(n \log n)$
(C) $O(n)$
(D) $O(n^2)$

142. What is the maximum flow problem in a graph concerned with?
(A) Finding the path with the maximum number of edges.
(B) Finding the maximum weight edge in the graph.
(C) Finding the maximum number of vertices in a connected component.
(D) Finding the maximum amount of flow that can be sent from a source to a sink vertex.

143. An NP-hard problem is a problem that is:
(A) At least as hard as the hardest problems in NP.
(B) Verifiable in polynomial time.
(C) Solvable in polynomial time.
(D) A problem that cannot be solved by a computer.

144. Which of the following is an advantage of a linked list over an array?
(A) Random access to elements is faster.
(B) Dynamic size; insertions and deletions are efficient.
(C) It requires less memory per element.
(D) Memory locality is better, leading to faster cache performance.

145. Arrange the following Big-O complexities from fastest growing to slowest growing.
1. $O(n!)$
2. $O(2^n)$
3. $O(n^2)$
4. $O(n \log n)$
(A) 4, 3, 2, 1
(B) 1, 2, 3, 4
(C) 1, 3, 2, 4
(D) 4, 2, 3, 1

146. A complete binary tree is a binary tree in which:
(A) All levels are completely filled except possibly the last level, which is filled from left to right.
(B) Every node has two children.
(C) The left subtree is always larger than the right subtree.
(D) All leaf nodes are at the same level.

Click Here to Buy book on Amazon

(B) P-3, Q-4, R-1, S-2
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-1, R-4, S-3

90. Consider the following statements.
Statement I : A B* tree is a B-tree variant that requires non-root nodes to be at least two-thirds full.
Statement II : The main benefit of a B* tree over a B-tree is improved storage utilization.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

91. The P vs NP problem is a major unsolved problem in computer science. It asks whether:
(A) every problem that can be solved can also be solved quickly.
(B) parallel computers are fundamentally faster than sequential computers.
(C) randomized algorithms are more powerful than deterministic algorithms.
(D) every problem whose solution can be quickly verified can also be quickly solved.

92. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The time complexity to find the height of a skewed binary tree with n nodes is O(n).
Reason (R) : In a skewed binary tree, the tree structure resembles a linked list, requiring traversal of all nodes to find the deepest leaf.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

93. The Rabin-Karp string matching algorithm uses which technique to efficiently compare substrings?
(A) Finite automata
(B) Suffix trees
(C) Hashing
(D) Dynamic programming

94. What is the key idea behind parallel algorithms for sorting?
(A) To use a faster comparison operation.
(B) To use a sorting network that is hard-wired for a specific number of inputs.
(C) To reduce the problem to a string matching problem.
(D) To perform multiple comparisons and swaps simultaneously on different processors.

95. A Las Vegas algorithm is a type of randomized algorithm that:
(A) always produces a correct answer, but its running time has a random element.
(B) has a deterministic running time but might produce an incorrect answer.
(C) is used for problems that have no optimal solution.
(D) is another name for an approximation algorithm.

96. Which data structure would be most efficient to find the 10th largest number out of a stream of one billion numbers?
(A) A sorted array
(B) A min-heap of size 10
(C) A max-heap of size 10
(D) A hash table

97. Consider the following statements about the Divide and Conquer paradigm.
I. It breaks a problem into several smaller, independent subproblems.
II. The subproblems are of the same type as the original problem.
III. The solutions to the subproblems are combined to form the solution to the original problem.
IV. Binary Search and Merge Sort are classic examples of this paradigm.
Which of the above statement(s) is/are correct?
(A) I and III only
(B) I, II, and IV only
(C) II, III, and IV only
(D) I, II, III, and IV

98. The task of merging two sorted lists of size 'm' and 'n' into a single sorted list takes how much time?
(A) $O(m+n)$
(B) $O(m \times n)$

NODIA

Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-3, Q-4, R-2, S-1
(C) P-4, Q-3, R-2, S-1
(D) P-1, Q-2, R-3, S-4

157. Which of the following statements about Kruskal's algorithm for MST is INCORRECT?
(A) It sorts all the edges in non-decreasing order of their weights.
(B) It adds edges to the MST one by one, ensuring no cycles are formed.
(C) It builds the MST from a single starting vertex, growing outwards.
(D) It is a greedy algorithm.

158. The concept of reducibility is fundamental to proving a problem is:
(A) In P
(B) Solvable by a greedy algorithm
(C) Having a logarithmic time complexity
(D) NP-Hard

159. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : In hashing, collisions are undesirable.
Reason (R) : Collisions increase the time required to access or insert an element, degrading the performance from O(1) to something worse.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

160. What is a key difference between Branch and Bound and Backtracking?
(A) Backtracking is for optimization problems while Branch and Bound is for decision problems.
(B) Branch and Bound uses a bounding function to prune the search space, while backtracking uses a simpler is promising check.
(C) Backtracking uses a queue while Branch and Bound uses a stack.
(D) There is no fundamental difference between them.

161. A linear search on an array of 'n' elements has a worst-case time complexity of:
(A) $O(1)$
(B) $O(n)$
(C) $O(\log n)$
(D) $O(n^2)$

NODIA

162. A binary tree that is not empty has one node with no parent. This node is called the:
(A) Root
(B) Leaf
(C) Sibling
(D) Ancestor

163. Which of the following are examples of Divide and Conquer algorithms?
I. Binary Search
II. Merge Sort
III. Quick Sort
IV. Strassen's Matrix Multiplication
V. Prim's Algorithm
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, and V only
(D) All are correct

164. The problem of determining if a graph has a Hamiltonian cycle is:
(A) In P
(B) Solvable using Dijkstra's algorithm
(C) A minimum spanning tree problem
(D) NP-complete

165. For a given linked list p, the C code `p->next = p->next->next;` is intended to:
(A) Insert a node after p
(B) Delete the node p itself
(C) Delete the node after p
(D) Reverse the link of p

166. A priority queue can be efficiently implemented using which data structure?
(A) A simple array
(B) A heap
(C) A linked list
(D) A stack

167. Which of the following is NOT a valid rotation in an AVL tree?
(A) Left-Right (LR) Rotation
(B) Right-Left (RL) Rotation
(C) Left-Left (LL) Rotation
(D) Top-Bottom (TB) Rotation

168. Consider the following statements.
Statement I : The main objective of algorithm design is to produce a correct and efficient solution to a problem.
Statement II : Time complexity and space complexity are the two main measures of an algorithm's efficiency.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

169. What is the role of the comparison tree in lower bound theory?
(A) To represent the structure of a B-tree.
(B) To visualize the state space of a backtracking algorithm.
(C) To show the parent-child relationships in a heap.
(D) To model the execution of any comparison-based algorithm.

170. Match List I with List II.
List I: (Graph Algorithm)
P. Breadth-First Search (BFS)
Q. Depth-First Search (DFS)
R. Dijkstra's Algorithm
S. Kruskal's Algorithm
List II: (Underlying Data Structure)
1. Stack (or recursion)
2. Queue
3. Disjoint Set Union (for cycle detection)
4. Priority Queue
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-4, R-1, S-3
(D) P-4, Q-1, R-2, S-3

171. If a problem is in NP, what does this guarantee?
(A) An efficient algorithm exists to solve it.
(B) It is impossible to solve.
(C) It requires exponential memory.
(D) A proposed solution can be checked for correctness in polynomial time.

172. The process of multiplying two polynomials can be made faster using which advanced algorithm?
(A) The Sieve of Eratosthenes
(B) Fast Fourier Transform (FFT)
(C) The Ford-Fulkerson method
(D) The Knuth-Morris-Pratt algorithm

173. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A B+ tree is often preferred over a B-tree for database indexing.
Reason (R) : The linked list structure of the leaf nodes in a B+ tree allows for efficient sequential processing of records (range queries).
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

174. Which of the following search algorithms requires the data to be sorted?
(A) Linear Search
(B) Breadth-First Search
(C) Depth-First Search
(D) Binary Search

175. The complexity of Strassen's matrix multiplication algorithm, a Divide and Conquer approach, is approximately:
(A) $O(n^{2.81})$
(B) $O(n^3)$
(C) $O(n^2 \log n)$
(D) $O(n^2)$

176. Which of the following statements about graph algorithms are true?
I. Dijkstra's algorithm is a greedy algorithm.
II. BFS finds the shortest path in terms of the number of edges.
III. DFS can be used for topological sorting.
IV. Prim's algorithm always starts from a single vertex.
Which of the above statement(s) is/are correct?
(A) I and II only

Downloaded from www.pdf.tube

(B) I, II, and IV only
(C) II, III, and IV only
(D) I, II, III, and IV

177. The height of a binary search tree with 'n' nodes in the worst case is:
(A) $O(n)$
(B) $O(\log n)$
(C) $O(n \log n)$
(D) $O(1)$

178. A randomized algorithm is one that:
(A) uses a random number at least once during its execution to make a decision.
(B) produces a random output.
(C) only works on randomized data.
(D) has a running time that is unpredictable.

179. A Deque (Double-ended queue) is a data structure where elements can be added or removed from:
(A) Only the front
(B) Only the rear
(C) Both the front and the rear
(D) Only the middle

180. A graph where edge weights can be negative, but there are no cycles of negative total weight, can have its single-source shortest paths found using:
(A) Prim's Algorithm
(B) Dijkstra's Algorithm
(C) Breadth-First Search
(D) Bellman-Ford Algorithm

181. Consider the following concepts of performance analysis.
I. Time complexity measures the runtime of an algorithm.
II. Space complexity measures the memory usage of an algorithm.
III. Asymptotic notations describe the behavior for large input sizes.
IV. Big-O provides an upper bound on the growth rate.
V. Big-Ω provides a lower bound on the growth rate.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, IV, and V only
(D) I, II, III, IV, and V

182. Which of the following is the most significant drawback of implementing a queue using a simple array?
(A) It is not possible to implement a queue using an array.
(B) The size is fixed, and it can lead to wasted space or overflow.
(C) Accessing elements is slow.
(D) Deletion is a very complex operation.

NODIA

183. Consider the following statements.
Statement I : The set of problems solvable in polynomial time by a deterministic Turing machine is called P.
Statement II : The set of problems solvable in polynomial time by a non-deterministic Turing machine is called NP.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

184. The solution to the recurrence T(n) = 8T(n/2) + $O(n^2)$ using the Master Theorem is:
(A) $O(n^3)$
(B) $O(n^2)$
(C) $O(n^2 \log n)$
(D) $O(n^{\log_2 8})$ which is $O(n^3)$

185. What is the purpose of an approximation algorithm?
(A) To find an exact solution to a problem very quickly.
(B) To approximate the time complexity of an algorithm.
(C) To find a near-optimal solution for an NP-hard problem in polynomial time.
(D) To use randomization to improve performance.

186. The number of edges in a complete graph with 'n' vertices is:
(A) $n\frac{n-1}{2}$
(B) $n-1$
(C) $n\frac{n+1}{2}$
(D) n

blocks for efficient algorithm design. Arrays offer a contiguous block of memory for storing elements of the same type, enabling constant-time access based on an index. This property makes them highly efficient for applications requiring frequent lookups. A specialized application of arrays is the representation of a Sparse Matrix, where most elements are zero; instead of storing the entire matrix, one can use arrays to store only the non-zero elements along with their indices, thus saving significant space. In contrast to the static or fixed-size nature of basic arrays, other linear structures offer dynamic behavior. Stacks operate on a Last-In, First-Out (LIFO) principle, making them suitable for tasks like function call management and expression evaluation. Queues, conversely, follow a First-In, First-Out (FIFO) principle, ideal for managing tasks in sequence, such as print jobs. A variation, the Priority Queue, allows elements to be dequeued based on an associated priority rather than their arrival order. For fully dynamic memory management, Linked Lists are employed. They consist of nodes containing data and a pointer to the next node, allowing for flexible insertion and deletion without the need to shift elements, a common overhead in array-based lists.

Based on the above source, attempt the following five questions.

251. According to the passage, what is the primary advantage of using arrays to represent a Sparse Matrix?
(A) It increases the speed of matrix multiplication.
(B) It conserves memory by storing only non-zero elements.
(C) It allows for the storage of heterogeneous data types.
(D) It simplifies the process of transposing the matrix.

252. Which data structure operates on a Last-In, First-Out (LIFO) principle?
(A) Stacks
(B) Queues
(C) Priority Queues
(D) Linked Lists

253. How does a Priority Queue differ from a standard Queue?
(A) A Priority Queue follows a LIFO principle, while a standard Queue follows FIFO.
(B) A Priority Queue can only store numerical data, unlike a standard Queue.
(C) A Priority Queue processes elements based on priority, not just arrival time.
(D) A Priority Queue is a static data structure, while a standard Queue is dynamic.

254. What characteristic of Linked Lists makes them particularly suitable for scenarios requiring frequent insertions and deletions?
(A) They provide constant-time access to elements via an index.
(B) They store elements in a contiguous block of memory.
(C) Their nodes are stored in a pre-sorted order.
(D) Their structure allows for element manipulation without shifting other elements.

255. A system is designed to manage a list of print jobs, ensuring that the first job submitted is the first one to be printed. Based on the principles described, which data structure would be most appropriate for this task?
(A) A Stack, because it manages tasks recursively.
(B) A Sparse Matrix, because it is efficient for storing job details.
(C) A Queue, because it processes elements in the order they are received.
(D) A Priority Queue, because some print jobs may be more urgent than others.

Direction : Please read the following source carefully.

Tree-based data structures provide a hierarchical way to organize data, enabling efficient searching and sorting operations. A basic Binary Tree is a structure where each node has at most two children. To improve traversal efficiency, a Threaded Binary Tree can be used, where null pointers are replaced with pointers (threads) to the in-order successor or predecessor, facilitating traversal without recursion or a stack. For ordered data, the Binary Search Tree (BST) maintains a specific property: for any given node, all values in its left subtree are less than the node's value, and all values in its right subtree are greater. While BSTs are efficient on average, they can degenerate into a

Click Here to Buy book on Amazon

(C) The ability to reduce a problem's time complexity.
(D) The ability to simplify the code of a program.

196. The time complexity of merging two sorted arrays of size n each into a single sorted array is:
(A) $O(n \log n)$
(B) $O(n^2)$
(C) $O(\log n)$
(D) $O(n)$

197. Which of the following statements about a simple linked list is true?
(A) It requires a contiguous block of memory.
(B) Insertion and deletion at the end are O(1) operations.
(C) It is a non-linear data structure.
(D) To access the i-th element, one must traverse the first i-1 elements.

198. The Floyd-Warshall algorithm is used to solve which problem?
(A) Single-source shortest paths.
(B) All-pairs shortest paths in a weighted graph.
(C) Minimum spanning tree.
(D) Maximum flow.

199. Consider the following statements.
Statement I : Heap sort is an in-place sorting algorithm.
Statement II : Heap sort is not a stable sorting algorithm.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

200. What is a collision in the context of hashing?
(A) When two different keys produce the same hash value.
(B) When the hash table becomes full.
(C) When a key is not found in the hash table.
(D) When the hash function takes too long to compute.

201. The number of null pointers in a non-empty binary tree with 'n' nodes is:
(A) n
(B) n - 1
(C) n + 1
(D) 2n

202. Which of the following is an example of a problem for which an approximation algorithm is often used?
(A) The Traveling Salesperson Problem (TSP)
(B) Sorting an array
(C) Searching in a binary search tree
(D) Calculating Fibonacci numbers

203. Consider the characteristics of greedy algorithms.
I. They make a sequence of choices.
II. For each decision point, they make the choice that seems best at the moment.
III. This local-optimal choice is never reconsidered.
IV. They are guaranteed to find the global optimum for all optimization problems.
V. Dijkstra's algorithm and Kruskal's algorithm are examples.
Which of the above statement(s) is/are correct?
(A) I, II, and III only
(B) I, II, III, and V only
(C) II, III, IV, and V only
(D) All are correct

NODIA

204. Which of the following best describes time complexity?
(A) The actual time in seconds an algorithm takes to run.
(B) A measure of how the running time of an algorithm scales with the size of the input.
(C) The number of lines of code in the algorithm.
(D) The amount of memory the algorithm uses.

205. If a hashing function is `h(key) = key % 7` and chaining is used for collision resolution, what will the hash table look like after inserting the keys 10, 25, 17, 3, 24?
(A) Index 3 will have a chain: 10 -> 24 -> 3
(B) Index 3 will have a chain: 10 -> 3 -> 17
(C) Index 3 will have a chain: 3 -> 17 -> 24
(D) Index 3 will have a chain: 3 -> 10. Index 4 will have a chain: 25 -> 17.

206. Match List I with List II.
List I: (Recurrence Relation)
P. $T(n) = 2T(n/2) + O(n)$
Q. $T(n) = T(n-1) + O(n)$
R. $T(n) = 2T(n/2) + O(1)$

S. $T(n) = T(n/2) + O(1)$

List II: (Time Complexity)

1. $O(n^2)$
2. $O(\log n)$
3. $O(n)$
4. $O(n \log n)$

Choose the correct option from those given below:

(A) P-4, Q-1, R-3, S-2
(B) P-4, Q-2, R-3, S-1
(C) P-1, Q-4, R-3, S-2
(D) P-3, Q-1, R-4, S-2

207. A graph algorithm that systematically explores the edges of a graph to visit every reachable vertex from a source vertex, exploring as far as possible along each branch before backtracking, is:

(A) Depth-First Search (DFS)
(B) Breadth-First Search (BFS)
(C) Prim's Algorithm
(D) Dijkstra's Algorithm

208. The main difference between a B-tree and a B+ tree is:

(A) B-trees are not balanced, while B+ trees are.
(B) B+ trees store data pointers only in leaf nodes, while B-trees can store them in internal nodes.
(C) B-trees can only have two children, while B+ trees can have many.
(D) B+ trees are used for in-memory structures, while B-trees are for disk-based structures.

209. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).

Assertion (A) : The QuickSort algorithm exhibits its worst-case performance when the input array is already sorted.

Reason (R) : In a sorted array, the pivot selection (e.g., always picking the last element) consistently results in highly unbalanced partitions, leading to a recurrence of T(n) = T(n-1) + O(n).

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

210. What is a major application of sparse matrices?

(A) Representing large graphs with few edges.
(B) Storing high-resolution photographic images.
(C) Implementing fundamental stacks and queues.
(D) Advanced number theoretic computations.

211. Which data structure is implicitly used in a recursive function call?

(A) Queue
(B) Linked List
(C) Stack
(D) Array

212. A B-tree of order 'm' has a maximum of how many children per node?

(A) m-1
(B) m+1
(C) 2m
(D) m

213. Consider the following statements.

Statement I : Randomized algorithms are probabilistic and may not produce the same result on every run for the same input.

Statement II : Approximation algorithms are deterministic and will always produce the same (approximate) result for the same input.

Which of the following is correct?

(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

214. The problem of finding the longest common subsequence of two strings is typically solved using:

(A) Greedy Algorithm
(B) Dynamic Programming
(C) Divide and Conquer
(D) Backtracking

215. Which of the following are generally true about parallel algorithms?

I. They aim to reduce the overall computation time.
II. They divide a problem into parts that can be solved concurrently.
III. Communication overhead between processors can be a significant performance

Downloaded from www.pdf.tube

factor.

IV. Not all problems can be effectively parallelized.

Which of the above statement(s) is/are correct?

(A) I and II only
(B) I, II, and III only
(C) II, III, and IV only
(D) I, II, III, and IV

216. The average case time complexity of Quick Sort is:

(A) $O(n \log n)$
(B) $O(n^2)$
(C) $O(n)$
(D) $O(\log n)$

217. In a threaded binary tree, if a node's right child pointer is NULL, it is made to point to the:

(A) In-order predecessor
(B) Pre-order successor
(C) In-order successor
(D) Root node

218. The time complexity of inserting a key into an AVL tree is:

(A) $O(\log n)$
(B) $O(1)$
(C) $O(n)$
(D) $O(n \log n)$

219. The greedy-choice property means:

(A) A locally optimal choice is never reconsidered.
(B) Every choice made must be the greediest in terms of resource consumption.
(C) The algorithm is greedy for memory.
(D) A globally optimal solution can be arrived at by making a sequence of locally optimal choices.

220. The final step in many Divide and Conquer algorithms is:

(A) Divide
(B) Combine
(C) Conquer
(D) Prune

221. Which of the following problems is NOT known to be NP-complete?

(A) 3-SAT
(B) Vertex Cover
(C) Finding the Minimum Spanning Tree
(D) Traveling Salesperson Problem (Decision Version)

222. A function $f(n) = O(g(n))$ implies that:

(A) $f(n)$ is asymptotically equal to $g(n)$.
(B) $f(n)$ is asymptotically greater than $g(n)$.
(C) $f(n)$ is asymptotically less than $g(n)$.
(D) $f(n)$ is asymptotically less than or equal to $g(n)$.

223. Identify the correct statements regarding specific algorithms.

I. The KMP algorithm is for string matching.
II. The FFT algorithm is for polynomial multiplication.
III. Euclid's algorithm is for finding the greatest common divisor.
IV. Sieve of Eratosthenes is a number theoretic algorithm for finding prime numbers.
V. Dijkstra's algorithm is a number theoretic algorithm.

Which of the above statement(s) is/are correct?

(A) I, II, and III only
(B) I, II, III, and IV only
(C) II, III, IV, and V only
(D) All are correct

224. The data structure used in a standard implementation of Breadth-First Search is:

(A) A stack
(B) A queue
(C) A priority queue
(D) A linked list

225. A sparse matrix is one where:

(A) The number of zero elements is much less than the number of non-zero elements.
(B) The number of zero elements is much greater than the number of non-zero elements.
(C) All elements are zero.
(D) All elements are non-zero.

226. Consider the following statements.

Statement I : In a B+ tree, all keys are present in the leaf nodes.

Statement II : In a B+ tree, internal nodes contain only keys to guide the search.

Which of the following is correct?

(A) Both Statement I and Statement II are true.

NODIA

(B) Both Statement I and Statement II are false.
(C) Statement I is true but Statement II is false.
(D) Statement I is false but Statement II is true.

227. The primary goal of a lower bound through reduction argument is to:
(A) Prove that problem A is easy by showing an easy problem B reduces to it.
(B) Reduce the memory requirement of an algorithm.
(C) Prove that problem A is hard by showing a known hard problem B reduces to it.
(D) Find a faster algorithm for a problem.

228. Which sorting algorithm works by repeatedly stepping through the list, comparing adjacent elements and swapping them if they are in the wrong order?
(A) Merge Sort
(B) Quick Sort
(C) Heap Sort
(D) Bubble Sort

229. If a graph traversal results in an ordering of vertices such that for every directed edge from vertex u to vertex v, u comes before v in the ordering, this ordering is known as a:
(A) Hamiltonian path
(B) Euler tour
(C) Topological sort
(D) Minimum spanning tree

230. The post-order traversal of a binary tree is Left-Right-Root. For the expression tree of `a + b * c`, what would be the post-order traversal?
(A) + a * b c
(B) a b c * +
(C) a + b * c
(D) a b + c *

231. The optimal substructure property, a key ingredient for dynamic programming, means that:
(A) The problem can be broken into smaller sub-problems.
(B) All sub-problems are independent of each other.
(C) The substructure of the problem is a tree.
(D) An optimal solution to the overall problem contains within it optimal solutions to sub-problems.

232. An array is a collection of items stored at:
(A) Contiguous memory locations
(B) Random memory locations
(C) A linked list of locations
(D) A tree of locations

233. The maximum number of edges in a simple graph with 'n' vertices is:
(A) n
(B) n/2
(C) n-1
(D) n(n-1)/2

234. In the context of parallel algorithms, a merge operation is often parallelized by:
(A) Using a rank-based approach to determine the final position of each element in parallel.
(B) Having one processor do all the work.
(C) Assigning elements to processors based on their value.
(D) Having each processor merge a small portion of the lists.

235. The concept of NP-completeness applies to which class of problems?
(A) Optimization problems
(B) Numerical problems
(C) All types of computational problems
(D) Decision problems

236. The worst-case height of a B-tree with 'n' keys and minimum degree 't' is proportional to:
(A) $O(n)$
(B) $O(\log_t n)$
(C) $O(n/t)$
(D) $O(\log_2 n)$

237. Which of the following is true for a complete graph with 'n' vertices?
(A) It has exactly 'n' edges.
(B) It is not connected.
(C) Every pair of distinct vertices is connected by a unique edge.
(D) It has a cycle of length 2.

238. The choice of the pivot element in Quick Sort is crucial for its performance. A poor pivot choice leads to:
(A) A more balanced partition.
(B) The best-case running time.
(C) The algorithm failing to sort correctly.

Click Here to Buy book on Amazon

(D) A highly unbalanced partition and worst-case performance.

239. What type of problem is the Maximum Flow problem?
(A) A graph optimization problem
(B) A sorting problem
(C) A string matching problem
(D) A number theoretic problem

240. The definition of $\Theta(g(n))$ (Big-Theta) is that a function $f(n)$ is bounded both above and below by constant multiples of $g(n)$ for large n. This represents a:
(A) Loose upper bound
(B) Tight asymptotic bound
(C) Loose lower bound
(D) Best-case scenario

241. Which of the following data structures allows for insertion and deletion of elements at both ends?
(A) Queue
(B) Deque
(C) Stack
(D) Priority Queue

242. In a comparison tree representing a sorting algorithm for 'n' items, the length of the longest path from the root to a leaf represents:
(A) The worst-case number of comparisons.
(B) The best-case number of comparisons.
(C) The average-case number of comparisons.
(D) The number of items 'n'.

243. The process of searching for an element in a binary search tree has a time complexity analogous to which other algorithm?
(A) Linear search
(B) Binary search
(C) Bubble sort
(D) Merge sort

244. A Las Vegas randomized algorithm for sorting (like randomized QuickSort) differs from a deterministic one in that:
(A) It might not sort the array correctly.
(B) It can only sort numbers, not strings.
(C) It has a higher worst-case complexity.
(D) Its running time is a random variable, but the output is always correct.

245. When we say a problem is reducible to another, we are using a tool from which theory?
(A) Graph Theory
(B) Number Theory
(C) Lower Bound Theory
(D) Complexity Theory

246. For which of the following operations is a doubly linked list more efficient than a singly linked list?
(A) Searching for an element.
(B) Deleting a node whose location is given.
(C) Inserting a node at the beginning.
(D) Traversing the list from start to end.

247. The primary use case for B-Trees and their variants (B+, B*) is in:
(A) Implementing in-memory caches.
(B) Storing data for compilers.
(C) Computer graphics and geometric algorithms.
(D) Implementing file systems and databases.

248. The solution to the recurrence T(n) = T(n-1) + 1 is:
(A) O(n)
(B) O(log n)
(C) O(n log n)
(D) O(n^2)

249. Which of the following is NOT a common method for resolving hash collisions?
(A) Linear Probing
(B) Quadratic Probing
(C) Separate Chaining
(D) Binary Probing

250. The performance of an algorithm is said to be asymptotically optimal if:
(A) Its running time matches the proven lower bound for the problem for large inputs.
(B) It is the fastest possible solution across all possible input variations.
(C) It uses the least amount of memory space required for any implementation.
(D) It is written in the fewest lines of code possible for a given language.

Direction : Please read the following source carefully.

Fundamental data structures provide the building

Downloaded from www.pdf.tube

187. Which of the following is NOT a greedy choice for finding a minimum spanning tree?
(A) Kruskal's: pick the next smallest edge that doesn't form a cycle.
(B) Prim's: pick the smallest edge connecting a visited vertex to an unvisited one.
(C) All are greedy choices, but for different problems.
(D) Dijkstra's: pick the unvisited vertex with the smallest distance from the source.

188. Match List I with List II.
List I: (Concept)
P. Lower Bound Theory
Q. NP-Completeness
R. Amortized Analysis
S. Randomized Algorithm
List II: (Focus)
1. Analyzing a sequence of operations to find an average cost
2. Using probability to improve algorithm performance
3. Establishing the minimum complexity for a problem
4. Classifying the hardest problems in NP
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-2, S-1
(C) P-3, Q-4, R-1, S-2
(D) P-4, Q-3, R-1, S-2

189. The pre-order traversal of a binary tree visits the nodes in which order?
(A) Left, Right, Root
(B) Left, Root, Right
(C) Root, Left, Right
(D) Right, Left, Root

190. In a max-heap data structure, the value of any node is:
(A) Less than or equal to the value of its children.
(B) Equal to the sum of its children's values.
(C) Always an even number.
(D) Greater than or equal to the value of its children.

191. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The complexity of polynomial multiplication can be reduced from O(n^2) to O(n log n).
Reason (R) : This reduction is achieved by using the Fast Fourier Transform (FFT) algorithm, which transforms the polynomials into a point-value representation, performs point-wise multiplication, and transforms back.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

192. Which data structure is used to manage the activation records of function calls in most programming languages?
(A) A queue
(B) A stack
(C) A heap
(D) A hash table

NODIA

193. The problem of sorting 'n' numbers is in which complexity class?
(A) It is not known.
(B) P
(C) NP-complete
(D) NP-hard but not NP-complete

194. Consider the algorithm design technique of Backtracking.
I. It is a form of brute-force search with a pruning mechanism.
II. It is suitable for problems that involve finding a sequence of choices.
III. The N-Queens problem is a classic example.
IV. It uses a depth-first search approach to explore the solution space.
Which of the above statement(s) is/are correct?
(A) I and III only
(B) I, II, and IV only
(C) II, III, and IV only
(D) I, II, III, and IV

195. What does the term reducibility signify in complexity theory?
(A) The ability to reduce the memory usage of an algorithm.
(B) The ability to transform an instance of one problem (A) into an instance of another problem (B).

linear structure in the worst case. To prevent this, self-balancing trees like the AVL Tree are used. An AVL tree is a height-balanced BST that performs rotations to maintain a balance factor (difference in heights of left and right subtrees) of -1, 0, or 1 for every node, guaranteeing logarithmic time complexity for operations. For very large datasets that do not fit in main memory, B-Trees and their variants (B+ Tree, B Tree) are optimal. These are multi-way search trees that minimize disk I/O by having a high branching factor, thus keeping the tree's height low. A B+ Tree, for example, stores all data records only at the leaf nodes, which are linked sequentially to allow for efficient range queries.*

Based on the above source, attempt the following five questions.

256. What is the primary purpose of using threads in a Threaded Binary Tree?
(A) To balance the height of the tree after insertions.
(B) To facilitate tree traversal without using a stack or recursion.
(C) To store additional data within the tree's null pointers.
(D) To connect multiple binary trees into a forest.

257. Which property must be maintained for any given node in a Binary Search Tree?
(A) The node's value must be greater than all values in its left and right subtrees.
(B) The height of the left and right subtrees must differ by at most one.
(C) All leaf nodes must be at the same depth from the root.
(D) Values in the left subtree are smaller, and values in the right subtree are larger.

258. According to the passage, why are AVL trees used as an improvement over standard Binary Search Trees?
(A) They guarantee logarithmic time complexity by maintaining height balance.
(B) They allow nodes to have more than two children.
(C) They use threads to make in-order traversal faster.
(D) They store all data records only at the leaf level for faster access.

259. For which scenario are B-Trees and their variants like the B+ Tree considered optimal?
(A) When performing frequent in-memory sorting operations.
(B) When implementing recursive algorithms with deep function calls.
(C) When managing datasets that are too large to fit in main memory.
(D) When data needs to be organized according to a LIFO principle.

260. A key distinction of a B+ Tree mentioned in the text is that it:
(A) stores all data records exclusively in its leaf nodes.
(B) performs rotations to maintain a balance factor of zero.
(C) replaces all null pointers with threads to predecessors.
(D) degenerates into a linear structure in the worst-case scenario.

Direction : Please read the following source carefully.

The performance analysis of algorithms is crucial for understanding their efficiency and scalability. This analysis primarily revolves around two key metrics: Time Complexity, which quantifies the amount of time an algorithm takes to run as a function of the length of the input, and Space Complexity, which measures the amount of memory space required. Rather than measuring exact runtimes, which depend on hardware and other external factors, complexity analysis uses Asymptotic Notation to describe the growth rate of an algorithm's resource usage. The three most common notations are Big-O (O), which provides an upper bound on the growth rate; Big-Omega (Ω), which provides a lower bound; and Big-Theta (Θ), which provides a tight bound, indicating that the growth rate is bounded both from above and below by the same function. This approach allows for a standardized, machine-independent comparison of algorithms. For recursive algorithms, their complexity is often expressed using Recurrence Relations. A recurrence relation is an equation or inequality that describes a function in terms of its value on smaller inputs. Solving these relations, using methods like the master theorem or substitution, is a fundamental technique for determining the time complexity of

Click Here to Buy book on Amazon

(B) To ensure that the algorithm makes a locally optimal choice at every node.
(C) To define the initial set of candidates for the solution.
(D) To estimate whether a subtree can contain a solution better than the current best.

274. Based on the passage, which statement correctly distinguishes Branch and Bound from Backtracking?
(A) Branch and Bound uses a bounding function to prune the search space, which is not a feature of simple backtracking.
(B) Backtracking is used for optimization problems, whereas Branch and Bound is used for decision problems.
(C) Backtracking explores the state-space tree, whereas Branch and Bound uses a greedy approach.
(D) Branch and Bound guarantees a globally optimal solution, whereas Backtracking does not.

275. A programmer needs to find a solution by exploring a large search space. The programmer decides to abandon partial solutions as soon as they violate the problem's constraints to save time. Which technique is being employed?
(A) Greedy Algorithm
(B) Dynamic Programming
(C) Backtracking
(D) Divide and Conquer

Direction : Please read the following source carefully.

Graphs are a fundamental non-linear data structure used to model relationships between objects, consisting of vertices (or nodes) and edges connecting them. To understand the structure of a graph and to solve problems based on it, traversal algorithms are essential. The two most fundamental graph traversal algorithms are Breadth-First Search (BFS) and Depth-First Search (DFS). Breadth-First Search explores the graph layer by layer, starting from a source vertex. It first visits all the vertices adjacent to the source, then the unvisited vertices adjacent to those, and so on. This level-by-level exploration is typically managed using a queue data structure. Because of this exploratory pattern, BFS is naturally suited for finding the shortest path between two nodes in an unweighted graph. In contrast, Depth-First Search explores as far as possible along each branch before backtracking. It starts at a source vertex and explores a path to its deepest unvisited descendant, then backtracks to explore other branches. This deep-dive approach is often implemented recursively or iteratively using a stack. DFS is widely used in applications such as topological sorting, finding connected components in a graph, and solving puzzles that involve exploring all possible configurations, like mazes.

Based on the above source, attempt the following five questions.

276. Which data structure is typically used to manage the exploration of vertices in a Breadth-First Search (BFS)?
(A) A queue
(B) A stack
(C) A priority queue
(D) A binary search tree

277. What is the characteristic exploration pattern of Depth-First Search (DFS)?
(A) It explores the graph layer by layer, visiting all neighbors of a node first.
(B) It randomly selects the next vertex to visit from the set of unvisited vertices.
(C) It always visits the vertex with the lowest associated weight first.
(D) It explores as deeply as possible along one path before backtracking.

NODIA

278. For which of the following problems is Breadth-First Search naturally suited, according to the passage?
(A) Finding the longest path in a weighted graph.
(B) Finding the shortest path in an unweighted graph.
(C) Performing a topological sort of a directed acyclic graph.
(D) Identifying the articulation points in a graph.

279. How does the implementation of DFS typically manage the order of vertices to visit?
(A) By using a queue to maintain a FIFO order.
(B) By using a min-heap to prioritize vertices.
(C) By using recursion or a stack to maintain a

(C) To ensure that all sub-problems are completely independent.
(D) To exclusively use a top-down recursive approach for problem-solving.

268. The passage distinguishes Dynamic Programming from Divide and Conquer based on the nature of their sub-problems. What is this key difference?
(A) Dynamic Programming has overlapping sub-problems, while Divide and Conquer has independent ones.
(B) Divide and Conquer sub-problems are larger than those in Dynamic Programming.
(C) Dynamic Programming is a top-down approach, while Divide and Conquer is bottom-up.
(D) Divide and Conquer cannot be implemented recursively, unlike Dynamic Programming.

269. A problem has an optimal solution that can be built from the optimal solutions of its sub-problems. What is this property called?
(A) Overlapping sub-problems
(B) Memoization structure
(C) Optimal substructure
(D) Recursive decomposition

270. A programmer chooses a bottom-up implementation, solving all the smallest sub-problems first and using their results to solve progressively larger ones. Which design technique does this approach exemplify?
(A) The divide step of the Divide and Conquer strategy.
(B) A memoization structure.
(C) The conquer step of the Divide and Conquer strategy.
(D) A typical implementation of Dynamic Programming.

Direction : Please read the following source carefully.

Beyond divide-and-conquer and dynamic programming, other design techniques offer powerful strategies for tackling optimization and search problems. Greedy Algorithms build a solution piece by piece, always choosing the option that offers the most obvious and immediate benefit at each step. This approach aims for a locally optimal choice in the hope that it will lead to a globally optimal solution. While simple and often efficient, greedy algorithms do not guarantee a globally optimal solution for all problems. For problems requiring a systematic exploration of all possible solutions, Backtracking is a suitable technique. It incrementally builds candidates for the solution and abandons a candidate (backtracks) as soon as it determines that the candidate cannot possibly be completed to a valid solution. This forms a state-space tree, which is pruned whenever the algorithm determines a path is non-viable. Branch and Bound is an enhancement of this state-space search, primarily used for optimization problems. Like backtracking, it explores the state-space tree, but it computes a bound at each node to determine if the optimal solution can possibly lie in the subtree rooted at that node. If the bound indicates that the subtree cannot contain a solution better than the best one found so far, the entire subtree is pruned, often leading to a more efficient search than a simple backtracking approach.

Based on the above source, attempt the following five questions.

271. What is the fundamental principle of a Greedy Algorithm?
(A) To explore all possible solutions systematically before choosing one.
(B) To break a problem down into smaller, independent sub-problems.
(C) To make a locally optimal choice at each stage of the problem.
(D) To store the results of sub-problems to avoid re-computation.

272. How does the Backtracking technique operate when searching for a solution?
(A) It randomly selects a path and hopes it leads to a solution.
(B) It builds a solution incrementally and discards a path once it is deemed non-viable.
(C) It computes a bound to prune branches of the search space.
(D) It always finds the globally optimal solution by making locally optimal choices.

273. What is the primary function of the 'bound' in the Branch and Bound algorithm?
(A) To limit the depth of the recursion in the state-space tree.

Downloaded from www.pdf.tube

Hierarchical structures are categorized as non-linear because they organize data in parent-child relationships rather than sequences. A group of these independent structures forms a forest. Every valid non-empty structure must possess a single top-level starting point. Mathematically, a connected acyclic graph with a specific number of vertices always contains exactly one fewer edge than the total node count.

10. (A) Binary Search - O(log n)
Searching in a sorted collection using a divide-and-conquer strategy repeatedly halves the search space. This logarithmic behavior makes the process extremely fast even for very large datasets. Other common algorithms like bubble sort or insertion sort typically exhibit quadratic growth in their average cases, while merge sort maintains a consistent log-linear complexity regardless of the initial data arrangement or distribution.

11. (D) To minimize the number of collisions and ensure a uniform distribution of keys.
Effective mapping functions aim to spread keys evenly across all available slots in a storage table. When multiple keys map to the same location, additional steps are required to resolve the conflict, which slows down operations. A well-designed function keeps the average number of elements per slot low, maintaining the efficiency of search and insertion operations near constant.

12. (B) 2, 3, 1, 4
Basic hierarchical structures preceded specialized versions designed for efficiency. Self-balancing search structures were introduced later to ensure consistent performance. As data storage moved to external disks, multi-way trees were developed to minimize input/output operations. Refined versions of these multi-way structures, which optimize sequential access and range queries, represent the more recent advancements in indexing technology for large-scale database management systems.

13. (C) Storing the results of subproblems to avoid re-computation.
This algorithmic paradigm focuses on partitioning a complex problem into smaller, independent parts. Each part is solved individually, and the results are then merged to reach a final conclusion. Saving intermediate results for reuse is actually a characteristic of dynamic programming, not this specific strategy. This approach is highly effective for tasks like sorting, searching, and large-scale mathematical calculations.

14. (A) Both Statement I and Statement II are true.
Optimized problem-solving involves identifying recurring smaller tasks within a larger operation. By solving each unique sub-task once and storing the result, the overall computational effort is significantly reduced. Calculating sequence values sequentially is a perfect demonstration of this, as each value relies on previously determined results. This technique transforms exponential time processes into much more efficient linear or polynomial ones.

15. (B) P-3, Q-1, R-4, S-2
Merge sort partitions data into halves recursively. The 0/1 knapsack problem requires evaluating sub-decisions to find the best total value. Fractional knapsack problems allow taking parts of items, making them suitable for choosing the best value-to-weight ratio at each step. The N-Queens puzzle involves building a solution and retracting steps when conflicts are detected, which is the core of backtracking.

16. (B) It makes a locally optimal choice at each stage to find a global optimum.
This strategy focuses on making the best possible decision at the current moment without considering future consequences. It assumes that a series of small, optimal steps will lead to an overall best solution. While this approach is very fast and efficient for certain problems like finding minimum spanning trees, it does not guarantee an optimal result for all types of challenges.

17. (B) Backtracking
Solving the placement of pieces on a board involves an incremental search through all possible configurations. When the current path violates constraints, the algorithm reverts to the previous step and attempts a different path. This systematic exploration prunes the search tree by avoiding invalid states early. This method is highly effective for combinatorial problems where multiple constraints must be satisfied simultaneously.

LIFO order.
(D) By using an adjacency matrix for direct access.

280. A problem requires finding if a path exists between two nodes in a maze. The goal is to explore one path completely to its end to see if it's the exit before trying another path. Which traversal algorithm aligns with this strategy?
(A) Breadth-First Search, as it explores all adjacent paths simultaneously.
(B) A greedy algorithm, as it would choose the most promising direction at each step.
(C) A B-Tree traversal, as it is optimized for pathfinding.
(D) Depth-First Search, as it explores a single path to its conclusion before backtracking.

Direction : Please read the following source carefully.

Beyond simple traversal, graph algorithms provide solutions to a variety of critical optimization problems. One such category is finding the Shortest Paths between vertices in a graph, which is fundamental in applications like network routing and logistics. The goal is to find a path between two vertices such that the sum of the weights of the edges in the path is minimized. Another essential problem is finding a Minimum Spanning Tree (MST) for a connected, undirected, and weighted graph. An MST is a subgraph that connects all the vertices together with the minimum possible total edge weight, without forming any cycles. This is crucial in designing networks, like laying electrical cables or water pipelines, where the objective is to connect all points with minimum cost. A different class of problem is the Maximum Flow problem, which is defined on a flow network—a directed graph where each edge has a capacity. The problem is to find the maximum rate at which material can flow from a designated source vertex to a sink vertex without exceeding the capacity of any edge. This has direct applications in modeling traffic, network bandwidth, and supply chains. While both MST and Shortest Path problems deal with minimizing edge weights, their objectives are distinct: MST connects all vertices, whereas Shortest Path focuses on the connection between two specific vertices.

Based on the above source, attempt the following five questions.

281. What is the primary objective of a Minimum Spanning Tree (MST) algorithm?
(A) To connect all vertices in a graph with the minimum total edge weight.
(B) To find the shortest path between a source and a destination vertex.
(C) To find the maximum rate of flow from a source to a sink.
(D) To visit every vertex in a graph exactly once.

282. The Maximum Flow problem is concerned with which of the following?
(A) Minimizing the total weight of edges in a network.
(B) Determining the maximum rate of flow through a network with capacities.
(C) Finding the shortest path in a directed graph.
(D) Connecting all nodes without forming a cycle.

283. In which of the following applications would a Shortest Path algorithm be most relevant?
(A) Designing a low-cost telecommunication network to connect all cities in a region.
(B) Finding the most efficient route for a GPS navigation system.
(C) Determining the maximum number of cars that can travel through a road system.
(D) Creating a spanning subgraph with no cycles.

284. According to the passage, how does a Minimum Spanning Tree problem differ from a Shortest Path problem?
(A) MST connects all vertices, while Shortest Path connects two specific vertices.
(B) MST applies to directed graphs, while Shortest Path applies to undirected graphs.
(C) MST maximizes total edge weight, while Shortest Path minimizes it.
(D) MST allows cycles, while Shortest Path does not.

285. A city planner needs to design a pipeline network to connect several neighborhoods to a water reservoir. The goal is to use the minimum possible length of pipe to ensure every neighborhood is connected. Which graph problem does this scenario represent?
(A) Maximum Flow

Click Here to Buy book on Amazon

(A) To reduce the memory consumption of the algorithm.
(B) To simplify the logic and implementation of the algorithm.
(C) To find an approximate solution instead of an exact one.
(D) To decrease the total computation time by using multiple processors.

297. In what situation are Approximation Algorithms most suitably used?
(A) When the problem is computationally intractable and an exact solution is not feasible.
(B) When a problem can be easily solved in polynomial time.
(C) When an algorithm must guarantee the optimal solution, regardless of time.
(D) When the algorithm's logic needs to incorporate random choices.

298. What is the defining characteristic of a Randomized Algorithm?
(A) It guarantees to find the optimal solution for any problem.
(B) It incorporates an element of randomness into its logic.
(C) It breaks a problem down into smaller, independent sub-problems.
(D) It runs exclusively on parallel processing hardware.

299. An engineer is developing a solution for an NP-hard logistics problem. The system must provide a good, but not necessarily perfect, routing plan quickly. The solution's cost must be guaranteed to be no more than 1.5 times the absolute minimum cost. Which type of algorithm is being developed?
(A) A Parallel Algorithm
(B) An Approximation Algorithm
(C) A Randomized Algorithm
(D) A standard sequential algorithm

300. Based on the passage, what is a potential trade-off when using a Randomized Algorithm compared to a deterministic one?
(A) A lack of guaranteed correctness or variable runtime in exchange for simplicity and efficiency.
(B) Increased complexity in exchange for guaranteed optimality.
(C) Slower average performance in exchange for less memory usage.
(D) The requirement for parallel hardware in exchange for faster execution.

NODIA

cornerstone of this field and signal processing is the Fast Fourier Transform (FFT). The FFT is a highly efficient algorithm to compute the Discrete Fourier Transform (DFT) and its inverse, drastically reducing the complexity of polynomial multiplication and filtering operations from quadratic to nearly linear time. String Matching Algorithms are another crucial category, addressing the problem of finding occurrences of a specific pattern string within a longer text string. These algorithms are vital for applications like text editors, search engines, and bioinformatics, where searching for DNA sequences is a common task. The development of efficient algorithms in these selected topics has been instrumental in advancing fields ranging from computer security and scientific computing to information retrieval and computational biology, demonstrating the broad impact of specialized algorithm design.

Based on the above source, attempt the following five questions.

291. Which domain of algorithms is described as being fundamental to modern cryptography?
(A) String Matching Algorithms
(B) Polynomial Arithmetic
(C) Number Theoretic Algorithms
(D) Fast Fourier Transform

292. What is the primary problem addressed by String Matching Algorithms?
(A) Finding occurrences of a pattern within a larger text.
(B) Sorting a list of strings in lexicographical order.
(C) Compressing a text string to save storage space.
(D) Multiplying two polynomials represented as strings.

293. The Fast Fourier Transform (FFT) is identified as a highly efficient algorithm for computing what?
(A) The prime factors of a large integer.
(B) The shortest path in a graph.
(C) The optimal alignment of two strings.
(D) The Discrete Fourier Transform (DFT).

294. An algorithm designed to efficiently multiply two polynomials falls under which category mentioned in the passage?
(A) Number Theoretic Algorithms
(B) Polynomial Arithmetic
(C) String Matching Algorithms
(D) Graph Algorithms

295. A bioinformatician is searching for a specific gene sequence within a vast genome database. This task is a direct application of which algorithmic category?
(A) Polynomial Arithmetic
(B) Number Theoretic Algorithms
(C) String Matching Algorithms
(D) Fast Fourier Transform

Direction : Please read the following source carefully.

The field of advanced algorithms explores computational models and techniques beyond the standard sequential approach. Parallel Algorithms are designed to run on machines with multiple processors, dividing a problem into parts that can be solved concurrently. For fundamental tasks like sorting, searching, and merging, parallel algorithms aim to achieve a significant speedup over their sequential counterparts by leveraging concurrent processing capabilities. The goal is to reduce the overall computation time by distributing the workload. For problems that are computationally intractable, such as many NP-hard problems, Approximation Algorithms are often employed. Instead of seeking a perfect, optimal solution, which may take an infeasible amount of time, these algorithms aim to find a solution that is provably close to the optimal one within a reasonable time frame. The quality of the solution is guaranteed to be within a certain factor of the optimum. Another approach for tackling difficult problems is the use of Randomized Algorithms. These algorithms incorporate an element of randomness as part of their logic. By making random choices at certain points, they can achieve good performance on average for all inputs, often with greater simplicity and efficiency than deterministic algorithms, although they may not guarantee the correct result 100% of the time or might have a variable runtime.

Based on the above source, attempt the following five questions.

296. What is the primary goal of developing Parallel Algorithms for tasks like sorting and searching?

NODIA

(B) Shortest Path
(C) Minimum Spanning Tree
(D) Graph Traversal

Direction : Please read the following source carefully.

Complexity theory is a field dedicated to classifying computational problems according to their inherent difficulty. A fundamental distinction is made between the P and NP classes. The class P (Polynomial time) consists of decision problems that can be solved by a deterministic algorithm in polynomial time. These are generally considered to be computationally tractable or efficiently solvable. The class NP (Nondeterministic Polynomial time) comprises decision problems for which a given solution can be verified for correctness in polynomial time. While it is known that P is a subset of NP, it is an open question whether P equals NP. Within NP, the most difficult problems are classified as NP-complete. A problem is NP-complete if it is in NP and it is also NP-hard, meaning that every other problem in NP can be reduced to it in polynomial time. This concept of Reducibility is a key tool; if problem A can be reduced to problem B, a solution for B can be used to solve A. Proving that a problem is NP-complete is significant because if a polynomial-time algorithm were ever found for just one NP-complete problem, it would imply that polynomial-time algorithms exist for all problems in NP, thereby proving P=NP.

Based on the above source, attempt the following five questions.

286. Which of the following best defines the complexity class P?
(A) Problems for which a solution can be verified in polynomial time.
(B) Problems that can only be solved by non-deterministic algorithms.
(C) Problems that are considered computationally intractable.
(D) Problems that can be solved by a deterministic algorithm in polynomial time.

287. What is the defining characteristic of a problem in the class NP?
(A) It can be solved in polynomial time on a deterministic machine.
(B) It requires exponential time to solve on any machine.
(C) A proposed solution to the problem can be verified in polynomial time.
(D) It is impossible to find an efficient algorithm for it.

288. According to the passage, what two conditions must a problem satisfy to be classified as NP-complete?
(A) It must be in P, and every problem in P can be reduced to it.
(B) It must be in NP, and it must be NP-hard.
(C) It must be solvable in polynomial time, and its solution must be verifiable in polynomial time.
(D) It must be NP-hard, but it cannot be in the class NP.

289. What is the role of Reducibility in complexity theory as described in the text?
(A) It is a method for reducing the time complexity of an algorithm.
(B) It is a technique to prove that one problem can be transformed into another.
(C) It is a way to prove that a problem belongs to the class P.
(D) It is used to decrease the memory requirements of a problem.

290. What would be the major implication if a polynomial-time algorithm was discovered for a single NP-complete problem?
(A) It would mean that P = NP.
(B) It would prove that NP is a subset of P.
(C) It would have no impact on the P versus NP question.
(D) It would prove that no other NP-complete problems exist.

Direction : Please read the following source carefully.

Beyond general-purpose algorithms, several specialized domains require tailored algorithmic solutions. Number Theoretic Algorithms focus on problems involving integers, such as primality testing, integer factorization, and solving modular arithmetic equations, which are fundamental to modern cryptography. Another specialized area is Polynomial Arithmetic, which involves developing efficient algorithms for operations like addition, multiplication, and division of polynomials. A

Downloaded from www.pdf.tube

ANSWERS

1. (B) Stack
Linear data structures store elements in a sequence. This specific structure follows the last-in, first-out method, meaning the most recently added element is the first one removed. Operations like push and pop occur only at the top of the container. This behavior is fundamental for managing recursion, undo mechanisms, and parsing expressions in computer science applications.

2. (C) Statement I is true but Statement II is false.
Arrays use contiguous memory blocks, enabling direct access to any element using an index in constant time. This makes searching by index extremely efficient. However, standard arrays have a fixed size determined at the time of declaration. Resizing typically involves allocating a new memory block and copying elements, rather than dynamic adjustment of the existing block during runtime.

3. (A) P-3, Q-4, R-1, S-2
Stacks handle function calls by storing return addresses. Queues facilitate resource allocation where the first request is serviced first. Priority queues enable scheduling by processing tasks based on importance rather than arrival time. Sparse matrices efficiently represent large grids containing mostly zero values by storing only non-zero elements, which significantly reduces the memory requirements for large datasets in scientific computing.

4. (D) Queue
This linear structure operates on the first-in, first-out principle, where elements enter from one side and exit from the opposite side. The rear pointer manages additions, while the front pointer manages deletions. This mechanism ensures that the processing order matches the arrival order exactly. It is essential for buffering, process scheduling, and managing shared resources in computing environments.

5. (C) Any node can be accessed directly using an index in O(1) time.
Unlike arrays, these structures consist of nodes distributed across non-contiguous memory. Each node maintains a reference to the subsequent element. Accessing a specific position requires traversing the chain sequentially starting from the head. This linear traversal results in a time complexity proportional to the position of the element, making random access inefficient compared to array-based implementations in programming.

NODIA

6. (A) Both A and R are true and R is the correct explanation of A.
The efficiency of this sorting method depends heavily on the partition quality. When the chosen pivot consistently splits the data into the most unbalanced sets, such as one empty set and the remainder, the recursive depth becomes linear. This behavior leads to a quadratic growth in comparisons. This scenario typically occurs when sorting already sorted or reverse-sorted data.

7. (D) A sparse matrix representation.
Standard two-dimensional arrays reserve memory for every possible cell regardless of its content. In scenarios where connectivity is minimal, most cells contain zeroes. Using a specialized representation avoids storing these redundant null values, focusing only on active connections. This approach optimizes memory usage and enhances the performance of mathematical operations by skipping computations involving unnecessary zero-value elements within the grid.

8. (A) Elements are processed based on a priority value, not insertion order.
This abstract data type extends the standard queue by associating a ranking with every item. Instead of simple chronological order, the element with the highest rank is always extracted first. If multiple items share the same rank, they are typically handled in the order they arrived. This structure is vital for operating system scheduling and implementing Dijkstra's shortest path algorithm.

9. (D) I, II, III, and IV

algorithms designed using recursive approaches, such as those based on the divide-and-conquer paradigm.

Based on the above source, attempt the following five questions.

261. What does the Space Complexity of an algorithm measure?
(A) The number of lines of code in the algorithm.
(B) The execution time as a function of the input size.
(C) The number of recursive calls made during execution.
(D) The amount of memory space it requires to run.

262. What is the primary purpose of using Asymptotic Notation in algorithm analysis?
(A) To calculate the exact execution time of an algorithm on a specific machine.
(B) To provide a machine-independent way of describing an algorithm's growth rate.
(C) To determine the minimum possible memory an algorithm can use.
(D) To express an algorithm's logic in a formal mathematical language.

263. Which asymptotic notation is used to describe a tight bound on an algorithm's growth rate?
(A) Big-O (O)
(B) Big-Omega (Ω)
(C) Big-Theta (Θ)
(D) Little-o (o)

264. According to the passage, Recurrence Relations are typically used to express the complexity of which type of algorithms?
(A) Iterative algorithms
(B) Greedy algorithms
(C) Randomized algorithms
(D) Recursive algorithms

265. Based on the text, what can be inferred about analyzing an algorithm's performance?
(A) It focuses on the growth rate of resource usage relative to input size.
(B) The analysis is only concerned with the best-case scenario.
(C) The primary goal is to find an upper bound on its time complexity.
(D) It is highly dependent on the specific hardware used for execution.

Direction : Please read the following source carefully.

Effective algorithm design often relies on established techniques that provide a framework for solving complex problems. The Divide and Conquer strategy is a prominent example, which involves three steps: first, breaking down a problem into several smaller, independent sub-problems of the same type; second, solving these sub-problems recursively; and third, combining their solutions to create a solution for the original problem. This top-down approach is fundamental to many efficient algorithms. In contrast, Dynamic Programming is typically applied to problems where sub-problems are not independent but overlap. To avoid re-computing the same sub-problems multiple times, dynamic programming solves every sub-problem just once and stores its result in a table or memoization structure for future use. This technique can be implemented either top-down with memoization or, more commonly, bottom-up, by solving all smaller sub-problems first and gradually building up solutions to larger ones. The key characteristic distinguishing problems suitable for dynamic programming is the presence of both optimal substructure, where an optimal solution can be constructed from optimal solutions of its sub-problems, and overlapping sub-problems, which allows for performance gains through memoization.

Based on the above source, attempt the following five questions.

266. Which of the following is described as the first step in the Divide and Conquer strategy?
(A) Combining the solutions of sub-problems.
(B) Breaking the problem into smaller sub-problems.
(C) Solving the sub-problems recursively.
(D) Storing the results of sub-problems in a table.

267. What is the core reason for using Dynamic Programming, as mentioned in the passage?
(A) To solve problems that cannot be broken down into sub-problems.
(B) To avoid the redundant computation of overlapping sub-problems.

Click Here to Buy book on Amazon

18. (D) I, II, III, IV, and V
Standard graph algorithms provide reliable ways to explore connectivity. Layer-based exploration using a queue determines the shortest path in unweighted networks. Path-based exploration using a stack or recursion is ideal for identifying components and cycles. Both methods process every vertex and edge, resulting in linear complexity relative to the total number of elements when using efficient list-based storage representations.

19. (C) The in-order successor or in-order predecessor of the node
Standard binary trees often contain many empty pointers that consume memory without providing utility. By replacing these null references with links to the previous or next nodes in a specific traversal order, the structure becomes more efficient. This modification allows for non-recursive navigation through the tree and avoids the need for a separate stack during the traversal of the nodes.

20. (D) For any given node, all values in its left subtree are less than the node's value, and all values in its right subtree are greater.
This fundamental ordering rule ensures that data is organized logically for searching. At every level, the relationship between a parent and its descendants is strictly maintained. This property allows search operations to eliminate half of the remaining tree at each step, similar to binary search on a list. Consequently, finding, adding, or removing elements remains efficient in a balanced structure.

21. (A) Both A and R are true and R is the correct explanation of A.
Maintaining a nearly equal height across all branches prevents the tree from becoming overly deep and skewed. By enforcing a strict limit on the height difference between sibling subtrees, the structure guarantees logarithmic time complexity for all primary operations. Whenever an operation violates this balance, specific structural rotations are applied to restore the equilibrium without losing the search tree properties.

22. (B) Data pointers are stored only in leaf nodes, making sequential access much faster.
Consolidating all actual data records at the bottom level allows for the inclusion of sibling links between leaf nodes. These links enable fast, linear scanning of data without traversing back up to internal nodes. Furthermore, internal nodes can store more keys since they do not hold data, which increases the branching factor and reduces the overall depth of the structure.

NODIA

23. (D) A hash table or a balanced binary search tree
Managing collections of unique items requires efficient ways to check membership and combine groups. Specialized indexing structures allow for rapid lookup and insertion, which are the building blocks of group-based operations. Using these methods ensures that calculating the overlap or the combination of different groups remains computationally feasible even as the size of the individual collections grows very large.

24. (B) I, II and III only
Asymptotic notations provide a mathematical framework for describing how resource requirements change with input size. Upper bounds define the maximum possible growth, while lower bounds establish the minimum. Tight bounds occur when these two limits coincide. Space complexity specifically refers to memory utilization, whereas the time spent during execution is measured by a separate metric known as time complexity.

25. (C) $T(n) = 2T\left(\frac{n}{2}\right) + O(n)$
This mathematical expression describes a process where a task is divided into two equal parts. Each part is handled independently through recursion. After the sub-tasks are complete, an additional step proportional to the input size is required to combine the results into a final sorted order. This structure is a hallmark of efficient divide-and-conquer algorithms used in sorting.

26. (A) It does not use any bounding function to prune the search space.
Optimization strategies often rely on estimating the best possible outcome from a given state to avoid exploring unproductive paths. By calculating these limits, the algorithm can skip entire sections of the search tree that are guaranteed to yield worse results than the current best solution. This pruning is the defining feature

63. (A) $2^{\{h+1\}} - 1$
In a binary structure where every level is full, the number of nodes doubles at each descending step. Summing these values across all levels results in a geometric series. The total count is one less than the number of nodes that would be present in the next level. This formula provides the upper limit on how much data a tree of a certain depth can store.

64. (B) 2e
In a network where connections are mutual, every individual link connects two distinct points. When counting the number of connections at each point and summing them up, every single link is counted twice, once for each of its endpoints. This fundamental relationship is a basic principle in graph theory and applies to any simple undirected network regardless of its specific shape or size.

65. (A) I, II, III, and V only
Computational complexity categorizes problems based on the resources required to solve them. Problems solvable in polynomial time are a subset of those whose solutions can be verified in polynomial time. The most difficult verifiable problems are those that every other verifiable problem can be converted into. While these classes are distinct, the question of whether they are truly equal remains an unproven hypothesis.

66. (A) O(n)
Basic sorting methods can be optimized to detect if a list is already in the correct order. By performing a single pass and checking for any swaps, the algorithm can finish early if no changes are needed. This results in a very fast linear performance when the input is perfectly sorted, although the performance is much worse for unsorted or randomly distributed data.

67. (C) A connected graph
A network is described this way if it is possible to travel from any one point to any other point by following the existing links. This property ensures there are no isolated sections or unreachable nodes within the structure. While some networks might have direct links between every single pair, the basic requirement for this classification is simply the existence of at least one path between all points.

68. (D) The difference between the height of the left subtree and the height of the right subtree.
Maintaining structural equilibrium in a search tree involves tracking the relative depths of sibling branches. This specific value indicates whether a node is balanced or leaning more heavily to one side. By keeping this difference within a very narrow range at every node, the tree remains short and compact, ensuring that all search operations can be completed in logarithmic time.

69. (C) All of the above.
Using arrays to store numerical coefficients allows for straightforward mathematical operations. Simple addition and subtraction can be done in linear time. More complex tasks like multiplication can be significantly accelerated using advanced recursive techniques. Efficient evaluation methods also allow for finding the value of the expression at a specific point with a minimal number of steps, demonstrating the versatility of this storage approach.

70. (B) The MST is guaranteed to contain the shortest path between any two vertices in the graph.
A structure that connects all points in a network with the minimum total weight focuses on the global cost of connection. This does not necessarily mean that the path between any two specific points within that structure is the shortest possible path in the original network. The goal is overall efficiency of the tree, not the optimization of individual paths between specific pairs.

71. (A) P-3, Q-4, R-2, S-1
Some methods find the smallest item and place it correctly, while others build a sorted section one item at a time. Partitioning data into smaller groups is a common high-speed approach. Using a heap structure to repeatedly find the largest item provides a different efficient way to organize data. Each of these strategies represents a unique approach to the fundamental problem of data organization.

72. (A) Randomized Algorithms
Incorporating chance into a computational process can lead to faster average performance or simpler solutions for difficult problems. These methods use random decisions to avoid worst-case scenarios that might slow down deterministic approaches. While the exact path

Downloaded from www.pdf.tube

where standard multiplication methods would be prohibitively slow for large datasets.

36. (D) 1, 4, 3, 2
Sorting efficiency varies significantly in the worst-case scenario. Some methods maintain a consistent log-linear performance regardless of data order. Others, while often fast, can degrade to quadratic performance under specific conditions, like poor pivot selection. Basic methods that rely on simple swaps of adjacent elements consistently exhibit quadratic growth, making them the least efficient for large datasets in their least favorable configurations.

37. (C) Breadth-First Search (BFS) uses a stack for its implementation.
Searching a graph layer by layer requires a mechanism that processes discovered nodes in the exact order they were found. This behavior is provided by a first-in, first-out container. Using a last-in, first-out container would instead result in exploring deep into a single branch before moving to siblings, which is characteristic of a completely different type of traversal method used in graphs.

38. (B) Ford-Fulkerson algorithm
Determining the maximum throughput in a network involves finding paths with available capacity and incrementally increasing the total volume. This iterative process continues until no more paths from the starting point to the destination can accommodate more flow. It is a fundamental method in operations research and networking, distinct from algorithms designed for finding the shortest paths or connecting all points minimally.

39. (B) I, II, and III only
Refined versions of multi-way trees improve storage efficiency by enforcing stricter requirements on how full each node must be. Instead of splitting a node immediately when it becomes full, the system attempts to share keys with adjacent siblings. This redistribution delays the creation of new nodes, leading to a more compact structure that utilizes available space more effectively than the standard versions.

40. (D) Statement I is false but Statement II is true.
Logarithmic functions grow much more slowly than linear ones, meaning a process with logarithmic complexity is actually faster for large inputs. In contrast, adding a logarithmic factor to a linear function increases its growth rate, making it slower than a purely linear process. Understanding these relative growth rates is essential for selecting the most efficient algorithm for large-scale computational tasks.

NODIA

41. (C) The solution found is guaranteed to be within a factor of rho of the optimum.
When a problem is too difficult to solve perfectly in a reasonable amount of time, specialized methods provide a result that is close to the best possible one. This specific metric quantifies the maximum possible distance between the provided solution and the true mathematical optimum. It gives developers a reliable guarantee on the quality of the result when using these faster methods.

42. (A) P-3, Q-4, R-1, S-2
Minimum spanning trees are built by connecting nodes with the lowest cost. Shortest path problems for non-negative weights use a greedy approach to find distances from a single source. Comprehensive algorithms can find distances between all pairs of nodes in a network. In unweighted networks, a layer-by-layer search is the most efficient way to determine the minimum number of connections.

43. (D) has a deterministic running time but has a small probability of producing an incorrect answer.
This type of probabilistic method prioritizes speed and consistent performance. While it always finishes within a set time frame, there is a known chance that the final result might be wrong. By running the process multiple times, the probability of error can be reduced to an extremely low level, making it practical for many complex scientific and mathematical applications.

44. (B) Managing function calls in a recursive program.
Sequential storage containers are excellent for grouping related records or building more complex structures. They provide fast access to data and are widely used for fixed-size lists.

to signal that the container is at its maximum capacity.

82. (A) A function that estimates a lower bound (for minimization) or an upper bound (for maximization) on the solution cost within a subtree.
Optimization searches can be significantly accelerated by estimating the potential value of pursuing a specific path. If the best possible outcome from a certain branch is already worse than the best solution found so far, the algorithm can safely ignore that entire branch. This estimation process acts as a filter, focusing the search on only the most promising areas of the solution space.

83. (C) 2, 1, 3, 4
Searching began with simple sequential checks of every item. As data became organized, more efficient methods that divide the search space were developed. Later, researchers explored specialized techniques using different mathematical principles, like golden ratios or estimated positions based on value distribution. This progression shows a move from simple, general-purpose methods to more sophisticated techniques that exploit specific properties of the data being searched.

84. (D) O(log n)
This mathematical relation describes a process that halves the problem size at each step while doing a fixed amount of work. This is characteristic of searching in a sorted list or a balanced tree. Because the size of the task decreases so rapidly, the total number of steps grows very slowly as the input size increases, making it much more efficient than linear processes.

85. (C) The Sieve of Eratosthenes for finding prime numbers.
Certain procedures focus specifically on the properties and relationships of integers. Finding prime numbers by systematically removing multiples of smaller numbers is a classic example of this field. These methods are essential for modern computer security, which relies on the difficulty of mathematical tasks involving large integers. They are distinct from algorithms that manage data structures or solve graph-based problems through general search techniques.

86. (A) I, II, III, and IV only
Representing connections between nodes can be done using either a grid of every possible pair or a list of actual neighbors for each point. Grids provide fast checks for individual connections but consume more space for sparse networks. Lists are more compact when few connections exist but can be slower for specific queries. Both methods are versatile and can accommodate weighted links by storing values instead of simple indicators.

87. (D) Queue
Fairness in processing requests is best achieved by a structure that ensures the first arrival is the first one served. This behavior perfectly models real-world situations like customer lines or task scheduling in a shared environment. It maintains the chronological order of events, ensuring that every request is handled in the exact order it was received, preventing any item from being delayed indefinitely.

88. (A) Searching
Retrieving a specific item from a collection is a common and vital operation in computing. The efficiency of this process depends on how the data is organized. In unsorted lists, it might require checking every item, while in sorted or indexed structures, the target can be found much more quickly. This task is fundamental to database queries, file systems, and many everyday software applications.

89. (B) P-3, Q-4, R-1, S-2
Exploring deep into a network is ideal for finding cycles. Layered exploration is the standard way to find the shortest distance in unweighted paths. Solving complex nested problems is naturally handled by self-referential functions. Evaluating mathematical notations where operators follow their operands is efficiently managed by a container that handles items in the reverse order of their arrival, which is a stack.

90. (A) Both Statement I and Statement II are true.
Specialized multi-way trees improve performance by requiring nodes to maintain a higher level of occupancy. By redistributing data to siblings rather than splitting immediately, the structure remains more compact and has a higher branching factor. This leads to better utilization

Click Here to Buy book on Amazon

exponential processes to much more manageable polynomial time complexities in practice.

54. (B) 0/1 Knapsack Problem Algorithm
Some optimization tasks require considering multiple combinations of items that cannot be split. In these cases, making the best choice at each individual step does not guarantee the best overall result. Instead, the problem must be solved by evaluating all possible sub-decisions and their consequences. Other common tasks, like finding shortest paths or connecting network points, can be solved effectively using simple step-by-step optimization.

55. (D) O(log n)
In an organized hierarchical structure where elements are evenly distributed, each comparison allows the search to skip approximately half of the remaining nodes. This halving process ensures that the number of steps required to find any specific item grows very slowly as the total number of nodes increases. This efficiency is why balanced trees are preferred for high-performance searching and indexing applications.

56. (C) P-4, Q-1, R-2, S-3
Large-scale storage relies on multi-way trees for efficient indexing. Self-correcting structures ensure that search paths remain short. Stacks provide the necessary memory management for nested function execution. Networks of nodes and connections are represented using grids or lists of neighbors. Each of these specialized tools is selected based on the specific requirements of the data organization or the computational task.

57. (C) Traversal
Systematic exploration of every element in a structure is necessary for tasks like printing, copying, or performing calculations on all stored values. In hierarchies, this can be done in several different orders, such as visiting parents before or after their children. The goal is to ensure that every part of the structure is visited exactly once in a predictable and consistent manner.

58. (A) Both Statement I and Statement II are true.
Different versions of multi-way trees handle data storage in distinct ways. One variant keeps all actual records at the lowest level, using internal levels only for navigation. Another version allows records to reside at any level, potentially shortening the search for frequently accessed items. Both designs are balanced and effective for managing large datasets, but offer different performance characteristics for sequential and random access.

59. (A) Both A and R are true and R is the correct explanation of A.
Exhaustive searching through all possible combinations is often too slow for practical use. By identifying and abandoning partial solutions that cannot possibly lead to a valid final result, an algorithm can ignore massive sections of the search space. This early termination significantly reduces the total number of configurations that must be checked, making complex combinatorial problems solvable in a reasonable timeframe.

60. (B) floor(log2 n) + 1
Searching a sorted list by repeatedly dividing it in half is highly efficient. In the least favorable scenario, the process continues until the search space is reduced to a single element. The number of times a set of size n can be halved until it reaches one is the base-two logarithm of n. The final step ensures the target is either found or confirmed as missing.

61. (C) Big Theta
While some notations only describe the maximum or minimum possible growth of a function, this specific notation defines a range where the function is both bounded from above and below. It signifies that the algorithm's performance will always stay within a certain growth rate, providing the most precise description of its long-term behavior as the input size becomes very large.

62. (D) 3, 1, 4, 2
Divide and conquer is a classical approach used for many foundational algorithms. Greedy methods and backtracking followed as researchers developed techniques for optimization and complex search problems. Dynamic programming, which optimizes these searches by identifying overlapping sub-problems, represents a more sophisticated formalization that emerged later. This progression reflects the increasing complexity of the problems being addressed in computer science over time.

Downloaded from www.pdf.tube

100. (D) No particular order
Visiting the root of a tree only after all its descendants have been processed does not provide data in a sorted sequence in a standard search tree. While in-order traversal results in sorted values, this specific bottom-up approach is more suitable for tasks like deleting a tree or evaluating mathematical expressions where results of sub-operations must be known before the main operation can be completed.

101. (C) B# Tree
Multi-way trees are essential for managing large-scale data, and several variations exist to optimize different aspects of performance. Some focus on basic balancing, others on sequential access through linked leaf nodes, and some on increasing storage utilization by redistributing data between siblings. While these established variants are widely used in databases and file systems, some named variations do not exist as standard data structures in the field.

102. (B) $T(n) = 2T(n/2) + O(n)$
This expression represents the classic divide-and-conquer strategy where a task is split into two halves, each handled recursively. After the sub-tasks are complete, an additional linear-time step is required to merge or combine the results. This model accurately reflects the work performed in efficient sorting algorithms, leading to an overall log-linear growth in the total number of operations required for processing.

103. (B) I, II, III, and IV only
Self-correcting search trees maintain a nearly perfect balance by enforcing a strict height limit between sibling branches. If any operation causes an imbalance, the tree's structure is adjusted through specific rotations. This ensures that the depth remains very small relative to the number of nodes, guaranteeing very fast searching. However, the worst-case search time remains logarithmic, never becoming linear as it might in unorganized structures.

104. (A) To enable traversal without using recursion or a stack.
Standard methods for navigating a hierarchy require extra memory to keep track of the path. By repurposing empty links to point directly to the next or previous elements in a specific sequence, a tree can be explored using only its own internal pointers. This modification simplifies the traversal process and saves memory, which can be particularly useful in environments with limited resources.

105. (D) Statement I is false but Statement II is true.
Decisions that involve choosing entire items without the ability to take partial amounts cannot be solved by simply picking the best current option. These require a more comprehensive evaluation of all combinations. In contrast, if items can be divided into smaller portions, always selecting the highest value-to-weight ratio is a guaranteed way to reach the most optimal total result for the entire problem.

106. (A) P-3, Q-2, R-4, S-1
Computational complexity labels describe how performance changes as the input size grows. Constant time means the effort is fixed, while linear time means it scales directly with size. Quadratic time describes a much faster increase in work, typically seen in nested loops. Logarithmic time represents the slowest growth among these, characteristic of efficient search processes that repeatedly halve the remaining work.

107. (D) Stack
Converting mathematical expressions involves reordering operands and operators while respecting precedence rules. A temporary storage container is needed to hold operators until their corresponding operands are processed. Because the most recently seen operators with higher priority must be placed into the result before those seen earlier, a last-in, first-out structure is the most appropriate and efficient tool for this task.

108. (B) Bellman-Ford algorithm
Finding the shortest path in a network can be complicated by links with negative values. While most greedy methods fail in this scenario, an iterative approach that repeatedly refines distance estimates for all edges can correctly handle them. This method can also identify the presence of cycles that would make a shortest path impossible to define, providing a robust solution for more complex networking scenarios.

NODIA

taken by the computer may vary between runs, these techniques are often highly effective for tasks like primality testing, load balancing, and large-scale data analysis.

73. (A) Both Statement I and Statement II are true.
Converting a known difficult problem into a new problem using a fast translation process proves that the new problem is at least as hard as the original. If the new problem could be solved easily, then the known difficult one would also be easy, which contradicts its classification. This method is the primary way researchers demonstrate that certain challenges are fundamentally difficult to solve efficiently.

74. (A) 2L - 1
In a binary structure where every node has either zero or two children, there is a fixed mathematical relationship between the number of terminal points and the total size. Each internal node effectively splits a path into two, eventually leading to more endpoints. For every new leaf added, an internal node must also exist to create that branch, resulting in the total node count always being odd.

75. (A) Both A and R are true and R is the correct explanation of A.
Greedy methods for pathfinding assume that once a point is reached with the lowest current cost, that cost will never be improved. However, if some connections have negative values, a longer path with a large negative weight could eventually result in a lower total cost than the initial short path. This possibility invalidates the core assumption of the greedy choice, making the method unreliable.

76. (B) Karatsuba algorithm
Traditional multiplication of large numbers requires many individual digit products. By using a recursive approach that breaks the numbers into smaller parts and combines them using a clever mathematical trick, the total number of necessary multiplications is reduced. This strategy provides a significant speed improvement for very large numbers, which is essential for cryptography and high-precision scientific calculations where standard methods are inefficient.

77. (C) Queue
Exploring a network level by level requires a way to track which points have been discovered but not yet visited. A first-in, first-out container ensures that points discovered earlier are processed before those discovered later. This ensures that the exploration moves outward in a consistent, expanding wave from the starting point, which is necessary for finding the shortest path in an unweighted graph.

78. (D) I, II, III, and IV
Mapping data to specific slots in a table is a core technique for fast retrieval. Conflicts are inevitable when different data items target the same slot, requiring effective resolution strategies like linking items or searching for alternate spots. A perfect mapping is rare but ideal as it eliminates these conflicts. The primary goal is always to spread data evenly to keep search times as fast as possible.

79. (A) O(1)
In a structure where each element points to both its predecessor and its successor, removing an item only requires updating the links of its neighbors. If the exact location of the item is already known, this process involves a fixed number of pointer reassignments. This efficiency is a major advantage over other structures where finding the preceding element or shifting data would require much more work.

80. (B) Omega(n log n)
When an algorithm makes decisions based only on comparing pairs of items, there is a mathematical limit on how fast it can work. This limit is derived from the total number of possible ways the items can be arranged. To distinguish between every possible order, the algorithm must perform at least a logarithmic number of comparisons relative to the total number of permutations, leading to this lower bound.

81. (B) The queue is full.
In a fixed-size storage area used for a circular queue, two pointers track the front and rear. As elements are added, the rear pointer moves around the circle. When the space immediately following the rear is the same as the current front position, no more items can be added without overwriting existing data. This condition is used

Click Here to Buy book on Amazon

However, tracking active subroutines requires a system that handles nested returns, which is naturally managed by a last-in, first-out structure. This ensures that the most recent call is resolved before returning to the previous one.

45. (B) I, II, and IV only
Connected node structures offer significant flexibility in data management. Linking nodes in both directions enables easy navigation, while joining the ends creates a continuous loop. These designs allow for very fast modifications at any point if the location is known. Unlike arrays, these structures do not require elements to be side-by-side in memory, allowing them to grow dynamically.

46. (D) A B C D - * E / +
Converting mathematical notations involves rearranging operands and operators based on their precedence and grouping. Subtraction inside the parentheses is handled first, followed by the multiplication and division in the order they appear. The addition is the final operation performed. This resulting notation is easily evaluated by computer systems using a stack, as it eliminates the need for parentheses and complex rules.

47. (A) log2 N
In a perfectly distributed hierarchical structure, each level contains twice as many nodes as the level above it. This exponential growth in capacity means the total number of levels required to store a set of items is very small relative to the total count. This relationship is fundamentally logarithmic, ensuring that searching through the structure remains efficient even as the population increases.

48. (B) Both A and R are true but R is not the correct explanation of A.
Identifying whether a number is divisible by others is a decision problem with a verifiable solution. If a factor is provided, confirming it is simple arithmetic. While this confirms the number is not prime, the overall categorization of primality testing in complexity classes involves deeper mathematical properties. The ability to verify one aspect doesn't fully explain the complexity of the entire primality problem.

49. (B) Merge Sort
Sorting stability is the property that ensures items with identical keys maintain their original relative positions after the process is complete. This is particularly important when sorting data by multiple criteria sequentially. While many efficient algorithms do not naturally preserve this order due to long-range swaps, this specific divide-and-conquer method carefully merges sub-lists to keep the original order of equal elements intact.

50. (A) Traversing a linked list to find the last element.
This mathematical relation describes a process where one unit of work is done at each step, and the problem size decreases by exactly one for the next step. Moving through a sequential chain of nodes to reach the end follows this pattern exactly. Each node visit takes constant time, and the process repeats until the entire length of the chain is covered.

51. (C) Sorting
Organizing data according to a specific key or value is a fundamental task in computing. It makes subsequent operations, like searching or merging, significantly faster and more efficient. Various methods exist to achieve this, ranging from simple comparisons of adjacent items to sophisticated recursive strategies that partition and group data, each offering different trade-offs in terms of speed and memory usage.

52. (D) Index 8
In a table where conflicts are resolved by checking the next available slot, keys that map to the same starting point will cluster. Several keys ending in five will initially target the same position. As each subsequent key finds its primary spot occupied, it moves down the table until it finds an empty space. This sequence leads to the final placement at the first open index.

53. (B) I, II, and III only
This problem-solving technique is characterized by solving unique sub-problems once and using those results to build up to a final solution. It is highly effective when a problem can be broken into parts that overlap. Contrary to some assumptions, this method is designed to significantly improve efficiency, often reducing

NODIA

of storage space and potentially fewer levels, which is highly beneficial for managing very large datasets stored on external disks.

91. (D) every problem whose solution can be quickly verified can also be quickly solved.
One of the most famous open questions in science asks if the ability to check a provided solution quickly implies the ability to find that solution just as fast. Currently, many problems have fast verification methods but only slow, exhaustive search methods for finding the answer. If the answer to this question were positive, it would revolutionize fields like cryptography, logistics, and artificial intelligence.

92. (A) Both A and R are true and R is the correct explanation of A.
When a hierarchical structure loses its balance and nodes only branch in one direction, it effectively becomes a simple sequence. To find the deepest point in such a structure, one must visit every single node, as there are no alternative paths or shortcuts. This linear behavior defeats the purpose of the hierarchy and leads to much slower performance for all standard tree operations.

93. (C) Hashing
Identifying a pattern within a large text can be accelerated by comparing numerical representations of text segments rather than individual characters. This method calculates a value for the target pattern and then checks segments of the text. By using a sliding window technique to update the value efficiently, the algorithm can skip many character-by-character comparisons, significantly speeding up the search for multiple pattern matches.

94. (D) To perform multiple comparisons and swaps simultaneously on different processors.
Sorting large datasets can be accelerated by distributing parts of the task across multiple processing units. Each unit works on a segment of the data independently, and then the results are merged or coordinated. This parallel approach aims to reduce the total time taken by utilizing the collective power of multiple units, overcoming the physical limitations of a single processing stream for massive information.

95. (A) always produces a correct answer, but its running time has a random element.
This type of probabilistic method is designed to be reliable and accurate in its output. While the time it takes to finish can vary because of the random decisions it makes during execution, it will never give a wrong result. It is often used in situations where a deterministic algorithm might be very slow in the worst-case scenario, providing a much faster average-case performance.

96. (B) A min-heap of size 10
To find a small number of top elements in a massive stream, a compact structure that tracks the current smallest of the top items is most efficient. As new items arrive, they are only kept if they are larger than this current minimum. This keeps the storage requirements very low while ensuring that the final contents of the structure are the largest values seen in the entire stream.

97. (D) I, II, III, and IV
Solving complex tasks by repeatedly breaking them into smaller, similar sub-problems is a powerful strategy. Each small part is handled individually, and then the individual solutions are merged to form the final result. This method is the basis for many foundational algorithms used for searching, sorting, and advanced mathematical calculations, offering a clear and structured way to handle large-scale data processing efficiently.

98. (A) $O(m + n)$
Combining two already organized lists into a single one involves comparing the current leading items from each list and picking the smallest. This process continues until all items from both sources have been processed. Since each item is visited and moved exactly once, the total effort required is directly proportional to the combined total number of elements in both of the original lists.

99. (C) A collection of disjoint trees.
In graph theory, any network that contains no cycles is categorized based on its connectivity. A single connected acyclic network is a tree. If a network consists of multiple such parts that are not linked to each other, the entire collection is given this descriptive name. Each separate part maintains its own hierarchical structure without any paths that loop back on themselves.

Downloaded from www.pdf.tube

that makes this technique more efficient than exhaustive searching.

27. (A) P-3, Q-1, R-2, S-4
Problems solvable quickly by standard computers belong to a specific class. Another group includes problems where a proposed solution can be checked rapidly. The most difficult problems in this second group have the unique property that every other problem in the class can be translated into them. This translation process is a vital technique for comparing the difficulty of different challenges.

28. (B) Establishing the minimum amount of work any algorithm must perform to solve a given problem.
This theory focuses on the inherent difficulty of a computational task, independent of any specific implementation. It determines the mathematical floor for the number of operations required, such as comparisons in sorting. Understanding these limits helps researchers identify when an algorithm is asymptotically optimal, meaning it is impossible to create a fundamentally faster solution for that specific problem type.

29. (C) n!
Sorting a set of unique elements involves determining which of the many possible permutations represents the correct ordered sequence. Each leaf in a decision model represents one of these possible outcomes. Since there are a specific number of ways to arrange a set of items, the model must have at least that many endpoints to distinguish between every possible input arrangement accurately.

30. (B) I, II, III, and IV only
Modern computation utilizes various strategies to handle complex tasks. Some use chance to improve average performance, while others find near-perfect solutions for problems that are too hard to solve exactly. Distributing work across multiple units can speed up processing. However, methods relying on chance do not always guarantee a correct result in every single run, depending on their design.

31. (A) A is also in P class.
If a complex task can be converted into a simpler one using a fast translation process, and the simpler task is already known to be solvable quickly, then the original task must also be solvable quickly. This logical link is a fundamental principle in complexity theory, allowing researchers to categorize problems based on how they relate to known benchmarks of computational difficulty.

32. (B) A string matching algorithm
Finding a specific pattern within a larger body of text can be optimized by avoiding redundant comparisons. This particular method pre-analyzes the pattern to determine how much to shift when a mismatch occurs. By using previously gained information about the characters already matched, the search proceeds in a single pass, resulting in a highly efficient linear time complexity for the process.

33. (A) Both A and R are true and R is the correct explanation of A.
Balanced multi-way structures maintain efficiency by limiting the number of keys and children in each node. A node with a specific number of keys naturally creates one more partition than the key count, leading to a corresponding number of children. These strict limits ensure the structure remains balanced and that the height remains logarithmic, which is essential for large-scale data storage.

34. (A) Finding the Minimum Spanning Tree of a graph
Certain graph problems can be solved using efficient strategies that run in time proportional to a power of the input size. For example, connecting all points in a network with the lowest total cost is a task with several well-known fast solutions. In contrast, many other famous challenges, like finding the shortest tour through multiple cities, currently lack such efficient solutions.

35. (C) $O(n \log n)$
Multiplying large mathematical expressions traditionally requires a number of operations proportional to the square of their size. By transforming these expressions into a different representation using a specialized recursive technique, the required work is reduced significantly. This improvement is crucial for high-speed signal processing, data compression, and performing large-scale numerical simulations

NODIA

109. (A) Both A and R are true and R is the correct explanation of A.
Solving complex problems often involves repeating the same smaller tasks many times. By identifying these overlaps and saving the results of each unique sub-task, a computer can avoid a massive amount of redundant work. This optimization technique is a powerful way to transform tasks that would otherwise take an impossibly long time into ones that can be solved very quickly and efficiently.

110. (D) It is a matrix where most of the elements are zero.
Grids often contain a large amount of empty or null information. Storing every single cell in a standard way is wasteful of both memory and processing power. Specialized representations focus only on the significant, non-zero data points. This approach is essential in fields like scientific computing and network analysis, where very large datasets typically have only a tiny fraction of active connections.

111. (A) Ascending sorted order.
Navigating a search tree by visiting the left branch, then the root, and finally the right branch ensures that values are processed in their natural numerical sequence. This property is a direct result of the tree's ordering rules, where smaller values always reside to the left and larger ones to the right. It provides a simple and efficient way to retrieve all stored data in a sorted list.

112. (B) 4, 1, 2, 3
Solving a problem using an incremental search requires first defining the space of all possible configurations. The algorithm then attempts to build a solution one step at a time, checking if the current path remains valid. If a path is viable, the search continues deeper; otherwise, the algorithm reverts to the previous step and tries a different direction, systematically exploring only the promising options.

113. (B) O(n)
In the worst-case scenario, a search tree can become completely unbalanced, with each node having only one child. This structure resembles a simple list, where every new item must be added at the very end. Finding the correct insertion point then requires visiting every existing node in sequence, resulting in a time complexity that grows directly with the total number of items stored.

114. (A) NP-complete
Determining if a logical expression can ever be true is a foundational problem in computer science. It belongs to a special class where a solution can be checked quickly, but finding one may take a very long time. This problem is significant because it was the first to be proven as one of the hardest in its class, meaning any other problem in that class can be translated into it.

115. (D) I, II, III, and IV
This linear structure processes items in a strict first-arrival sequence. Adding elements to the back and removing them from the front ensures that the original order is preserved. This behavior is ideal for managing any situation where fairness and sequence are important, such as handling data packets in a network, scheduling print jobs, or modeling a physical line of people waiting for service.

116. (D) The list is partitioned, and each processor searches one partition.
When multiple processors are available, a large searching task can be divided by giving each unit a specific segment of the data to check. This simultaneous exploration can find the target much faster than a single processor checking every part in sequence. This strategy is a simple and effective way to utilize parallel hardware, although it requires careful coordination to manage the distribution and collection of results.

117. (C) 1.5
This value represents a guaranteed limit on the quality of a solution produced by a fast, non-optimal method. It indicates that the provided result will not exceed the best possible result by more than a specific percentage. In this case, the final tour length is at most fifty percent longer than the true minimum, providing a reliable bound for problems that are too difficult to solve perfectly.

118. (B) Polynomial Arithmetic
Manipulating large mathematical expressions can be significantly accelerated by converting them into a different form. Instead of traditional

Click Here to Buy book on Amazon

For an algorithm whose work increases with the square of the input size, doubling the data leads to a significant jump in the total number of operations. Mathematically, the new workload is two squared times the original. This rapid increase in effort is why such algorithms are generally avoided for very large datasets, as the time required to finish becomes much longer than for more efficient linear solutions.

155. (C) Statement I is true but Statement II is false.
While dividing a task among several processing units can speed up the process, it is rarely a perfect improvement. There is always some overhead involved in coordinating the different units and merging their results. Furthermore, some parts of a task might be impossible to parallelize, which limits the maximum possible speedup regardless of how many additional processing units are added to the system.

156. (B) P-3, Q-4, R-2, S-1
Organized hierarchies can take many forms to suit different needs. Some focus on simple numerical ordering for fast searching. Others add self-balancing rules to ensure consistent speed. Specialized versions store data only at the bottom for efficient scanning, while some add extra links to navigate without needing external memory. Each design offers a unique set of trade-offs for searching, inserting, and retrieving information.

157. (C) It builds the MST from a single starting vertex, growing outwards.
This optimization method for connecting points in a network takes a global view by considering all available links simultaneously. It systematically picks the smallest remaining links that do not create a loop, regardless of where they are in the network. This approach eventually merges separate connected groups into a single tree that covers all points with the lowest possible total weight across the entire network.

158. (D) NP-Hard
Demonstrating that a problem is at least as difficult as a whole class of known hard problems is done by showing how to transform any of those hard problems into the new one. This translation process is a key tool in complexity theory, allowing researchers to categorize the inherent difficulty of a challenge and predict whether an efficient solution is likely to exist.

159. (A) Both A and R are true and R is the correct explanation of A.
Mapping data to specific spots in a table works best when each item gets its own unique location. When multiple items target the same spot, the system must perform extra work to find a different place for them or to manage a list of items at that location. This extra work slows down both adding and finding information, reducing the overall efficiency of the data storage system.

160. (B) Branch and Bound uses a bounding function to prune the search space, while backtracking uses a simpler is promising check.
Both techniques involve searching through a space of possible solutions, but they differ in how they evaluate which paths to follow. One uses a simple check for validity at each step. The other uses more sophisticated estimates to predict the best possible outcome of a branch, allowing it to ignore large areas that are guaranteed to be worse than the current best solution, which is especially important for optimization.

161. (B) O(n)
Checking every item in a collection one by one is the most basic search method. In the worst-case scenario, the target is the very last item or is not present at all. In this case, the number of steps required is directly proportional to the total number of items in the list. While simple to implement, this method becomes very slow as the amount of data increases.

162. (A) Root
Every non-empty hierarchical structure begins at a single, unique starting point that has no parent above it. This point serves as the ultimate ancestor for every other node in the entire structure. All paths through the structure originate here, and it is the standard point of entry for operations like searching, traversing, or modifying the hierarchy as a whole in computer science.

163. (B) I, II, III, and IV only
Many foundational algorithms for searching, sorting, and complex mathematical tasks work by repeatedly splitting a large problem into smaller, independent parts. Each small part is

NODIA

Certain sorting methods are particularly efficient when the data is already mostly organized. By building the sorted section one item at a time and only moving items when they are out of place, the algorithm can complete the task with a very small number of comparisons and movements. In a perfectly sorted list, it only needs to check each item once, resulting in very fast performance.

128. (D) I, II, III, IV, and V
Hierarchical structures are incredibly versatile tools in computing. They range from simple binary trees to specialized versions that maintain balance, allow for many children, or include extra links for faster navigation. Each variant is designed to optimize specific operations like searching, sorting, or data retrieval. Together, they form the backbone of modern data management, enabling everything from simple file systems to complex, high-performance database indexing.

129. (B) Decision trees (Comparison trees)
The mathematical limit on the speed of sorting is established by modeling every possible sequence of comparisons as paths in a large diagram. Each endpoint of this diagram represents a unique ordering of the data. To account for every possible arrangement, the diagram must have a certain depth. This depth represents the minimum number of comparisons that any algorithm must perform to correctly identify the sorted order.

130. (A) Traversal
This operation is a fundamental task that involves visiting every single item in a collection once. It is a prerequisite for many other activities, such as displaying data, searching for a specific value, or applying a calculation to every element. Regardless of the internal organization of the data, a systematic way to move through all its components is essential for effectively managing and utilizing information.

131. (C) Internal nodes store both keys and data records.
A key feature of some multi-way trees is the separation of navigation and storage. While top levels are used to guide the search, the actual data is kept entirely at the bottom level. This allows the internal levels to hold more keys, leading to a wider branching factor and a shorter tree. This design also facilitates efficient sequential scanning of the data at the bottom level.

132. (D) O(V * E)
Finding the shortest path in a network that might have negative-value links requires a comprehensive approach. This algorithm works by repeatedly updating the distance estimates for every single connection in the network. It performs this update process multiple times, once for each point in the network. This ensures that the final distances are accurate even when paths involve complex combinations of positive and negative link values.

133. (A) I, II, and IV only
Decision problems are grouped into classes based on how quickly they can be solved or verified. Problems that can be solved quickly form one group, which is a subset of those that can be verified quickly. While many believe these two groups are different, no one has yet proven it. If they were actually the same, many currently difficult tasks would suddenly become very easy to solve.

134. (B) Backtracking
This problem-solving technique builds a solution piece by piece. At each step, it checks if the current configuration could potentially lead to a valid final result. If it determines that the current path is a dead end, it immediately discards it and returns to the previous decision point to try a different option. This selective exploration prevents the algorithm from wasting time on clearly incorrect paths.

NODIA

135. (D) DFS Forest/Tree
Exploring a network by following paths as deep as possible results in a structural representation of the nodes and the connections traversed. This result shows the order in which parts of the network were discovered and reveals the underlying hierarchy of the search process. It is a valuable tool for understanding connectivity, identifying cycles, and determining the overall structure of complex, interconnected datasets in computer science.

136. (A) P-3, Q-1, R-4, S-2
Identifying connected parts of a network is easily done by exploring all paths from a starting

multiplication, which is slow, this technique uses a recursive method to transform the data, perform a simpler operation, and then transform it back. This approach is fundamental for high-speed calculations in signals, images, and other areas where large-scale numerical data must be processed very rapidly.

119. (D) A Union-Find or Disjoint Set Union (DSU) data structure.
Managing groups of items that do not overlap requires a specialized way to track which group an item belongs to and to combine different groups. This specific tool provides two main operations: identifying the representative of a group and merging two separate groups into one. These operations are highly efficient and are used in algorithms for finding minimum spanning trees and tracking connected components.

120. (A) The size of the input.
The efficiency of an algorithm is mathematically analyzed based on how the required time and memory scale as the amount of data increases. This approach provides a universal measure that is independent of specific hardware, programming languages, or individual coding styles. It allows researchers to predict how a solution will behave as it moves from small test cases to massive real-world datasets in production.

121. (A) P-3, Q-1, R-4, S-2
Solving problems involves different strategic approaches. Some find solutions by exploring all possibilities and pruning invalid ones. Others optimize by reusing results of overlapping sub-problems. Partitioning a task into independent parts is another common method. Some simple optimization tasks can be solved by making the best possible choice at each step without ever looking back, which is a hallmark of the greedy approach.

122. (A) It is the amount of memory space required by an algorithm, as a function of the input size.
This metric quantifies the total memory needed by a computer program during its execution, including both the storage for the input data and any extra space required for calculations. Like time complexity, it is expressed as a function of the input size to show how the memory demands grow. It is a critical factor in designing software for devices with limited memory or handling massive datasets.

123. (C) Statement I is true but Statement II is false.
A specific mathematical theorem provides a simple way to determine the complexity of recursive algorithms that divide their work into equal parts. While it is very useful for many common algorithms, it cannot be applied to all recurrence relations. Some formulas do not fit the required structure or have specific terms that fall outside the theorem's defined cases, requiring more advanced analysis methods to solve.

124. (B) The Fibonacci numbers
The least number of nodes required to maintain a certain height in a self-balancing tree follows a specific mathematical pattern. Each level must have enough nodes to support the balance requirement, which relies on the heights of its sub-branches. This recursive dependency mirrors the structure of a well-known numerical sequence, showing the deep link between discrete mathematics and the design of efficient, balanced data storage structures.

125. (A) Both A and R are true and R is the correct explanation of A.
When a network has very few connections between its points, using a full grid to store every possible link is highly inefficient because most of the grid will be empty. Instead, a more compact representation that only records the actual links for each point saves a significant amount of memory. This approach is essential for handling massive real-world networks where the number of connections is relatively small.

126. (D) It is one of the hardest problems in the NP class.
Problems in this category share a unique and significant property: they are among the most difficult challenges whose solutions can be verified quickly. If an efficient way to solve any one of these problems were ever found, then every other problem that can be verified quickly could also be solved just as fast. They represent the frontier of what is currently solvable in reasonable time.

127. (A) Insertion Sort

Downloaded from www.pdf.tube

small inputs, while others remain practical for much larger datasets. Factorial and exponential growth represent the fastest-increasing categories, followed by quadratic and then log-linear growth. Understanding these relative speeds is vital for selecting the right approach to solve a problem based on the expected amount of data to be processed.

146. (A) All levels are completely filled except possibly the last level, which is filled from left to right.
This specific structural property ensures that a hierarchy is packed as tightly as possible. Every level above the bottom is full, and the bottom level is populated starting from one side. This organization allows the structure to be represented efficiently in memory using a simple sequence, which is why it is often used as the basis for high-performance priority-based data structures in computing.

147. (C) Dynamic Programming (top-down approach)
Optimization can be achieved by remembering the results of sub-tasks that have already been solved. When the same sub-task is encountered again, the computer simply retrieves the saved answer instead of recalculating it. This technique is a key part of more sophisticated problem-solving methods, allowing them to handle problems with overlapping parts much more efficiently than a simple recursive approach that would repeat the same work many times.

148. (C) I, II, III, IV, and V
Different ways of organizing data offer unique performance benefits. Some are simple and efficient for small sets, while others provide consistent speed for large amounts of information. Stability, memory usage, and the number of individual swaps are all important factors. Together, these properties define the character of each method, allowing developers to choose the most suitable one based on the specific constraints and requirements of their software application.

149. (C) Number Theoretic Algorithm
Finding the largest number that divides two integers without a remainder is a classic mathematical task. The method involves repeatedly replacing the larger number with the remainder of its division by the smaller one. This efficient procedure is a fundamental building block in mathematics and is essential for many modern applications, including the complex calculations used to secure communications and data on the internet today.

150. (B) Acyclic graph
Any network of connections where it is impossible to start at a point and return to it by following the links is given this name. These structures are essential for representing dependencies, hierarchies, and processes that move in a single direction. They are fundamentally different from networks that allow for loops, which can make pathfinding and other analysis tasks much more complex and computationally intensive to solve.

NODIA

151. (D) To describe the limiting behavior of a function when the argument tends towards a particular value or infinity.
Asymptotic notation provides a high-level mathematical way to compare the efficiency of different algorithms. It focuses on how the resource requirements grow as the amount of data becomes very large, ignoring minor details and constants. This allows researchers to focus on the fundamental efficiency of a solution, providing a clear way to predict how it will perform in large-scale real-world scenarios.

152. (A) 2t - 1
In a balanced multi-way structure, each node has a fixed capacity for storage. This capacity is defined by a parameter that determines the minimum and maximum number of entries allowed. These limits ensure the tree stays balanced and that each node is sufficiently full, which is vital for maintaining high performance during data retrieval and keeping the overall height of the structure small for large datasets.

153. (D) I, II, III, and IV
This linear structure is defined by its last-in, first-out behavior. It is essential for any task where the most recent item must be handled first, such as tracking active functions, reversing a sequence of characters, or managing an undo history. Adding and removing items at one end provides a simple and efficient way to manage data that follows this specific and common organizational pattern.

154. (A) 4

190. (D) Greater than or equal to the value of its children.
A specialized hierarchical structure can be organized so that the most significant item is always at the top. This is achieved by enforcing a rule that every node must be at least as large as any of its descendants. This simple rule ensures that the maximum value in the entire collection is immediately accessible, making it an ideal tool for tasks that require frequent access to the largest available element.

191. (A) Both A and R are true and R is the correct explanation of A.
Multiplying large mathematical expressions can be accelerated by transforming the data into a different numerical form. This transformation allows for simpler operations on the components. Once the operations are done, the data is transformed back. This sophisticated approach is much faster than traditional methods and is essential for high-speed processing in many modern scientific and engineering applications involving large-scale numerical calculations and signal analysis.

192. (B) A stack
Programming languages manage active subroutines by keeping track of where each one was called from. Because functions can call other functions, creating a nested structure, a container that handles the most recently added item first is needed. This ensure that when a function finishes, the computer can correctly return to the exact point in the previous function where it left off, maintaining the correct execution flow.

193. (B) P
Organizing data in a specific sequence is a fundamental task that can be solved very efficiently. There are many well-known methods for this that finish in a time proportional to a small power of the input size. Because these efficient solutions exist, the task belongs to the category of problems that are considered easy for modern computers to solve, even when dealing with very large amounts of information.

194. (D) I, II, III, and IV
This problem-solving technique explores all potential solutions by building them one piece at a time and discarding those that violate constraints. It is particularly effective for combinatorial tasks like placing items on a board without conflicts. By systematically trying different options and reverting when necessary, it ensures that all valid possibilities are considered while avoiding unnecessary work on paths that are guaranteed to fail.

195. (B) The ability to transform an instance of one problem (A) into an instance of another problem (B).
Comparing the difficulty of two different tasks is done by showing how to translate one into the other. If this translation process is fast, it implies that the first task is no harder than the second. This concept is a fundamental tool for categorizing problems into different groups based on how they relate to each other and to known benchmarks of computational difficulty in the field of science.

196. (D) O(n)
Combining two already organized lists into a single one is a very efficient process. It only requires looking at the next item from each list and picking the smallest. Every item is moved exactly once from its original list to the new, combined list. Because the total work is directly proportional to the total number of items, the process is very fast even for very large datasets.

197. (D) To access the i-th element, one must traverse the first i-1 elements.
In a chain of connected items where each one only knows about its immediate neighbor, finding a specific item requires starting at the beginning and moving through the chain one step at a time. This linear access is very different from other structures where any item can be reached instantly. This characteristic makes sequential searching much slower as the length of the chain increases over time.

198. (B) All-pairs shortest paths in a weighted graph.
Finding the minimum distance between every possible pair of points in a network is a comprehensive task. It involves an iterative process that considers every point as a potential intermediate step for every other pair. This ensures that the final result for every pair is the absolute shortest path possible, taking into

Click Here to Buy book on Amazon

handled individually, and the results are merged. This structured approach is highly effective for many tasks, but it is not the only way to solve problems; some optimization tasks are handled more efficiently by making a single sequence of best-choice decisions.

164. (D) NP-complete
Finding a path that visits every single point in a network exactly once and returns to the start is a very difficult task. It belongs to a class of problems where a provided answer is easy to check, but finding the answer is currently impossible to do quickly for large networks. This specific problem is a classic example used in the study of computational complexity and algorithm design.

165. (C) Delete the node after p
Modifying a chain of connected nodes involves changing where the links point. By making a node point directly to the one after its current neighbor, the middle neighbor is effectively bypassed and removed from the chain. This simple update is very efficient because it only requires changing a single link, demonstrating the power and flexibility of linked structures for managing dynamic lists of information.

166. (B) A heap
Managing a collection of items where the most important one is always processed first can be done efficiently using a specialized hierarchical structure. This organization ensures that the highest-priority item is always at the top and can be found in constant time. Adding new items or removing the top one takes a small amount of work that grows very slowly as the total number of items increases.

167. (D) Top-Bottom (TB) Rotation
Self-balancing trees maintain their equilibrium through four specific types of structural adjustments. These adjustments involve shifting nodes and their sub-branches to correct imbalances that occur when new data is added or removed. These four methods are the only standard ways to restore the required height constraints while keeping the tree's search properties intact; other types of adjustments are not part of the standard balancing procedures.

168. (A) Both Statement I and Statement II are true.
Effective software development relies on creating solutions that are both accurate and fast. The efficiency of these solutions is measured by how much time they take and how much memory they consume. These two factors are the primary benchmarks used by computer scientists to compare different approaches and to select the best one for a specific task based on the available resources and data size.

169. (D) To model the execution of any comparison-based algorithm.
Analyzing the limits of certain types of tasks involves creating a diagram that shows every possible decision an algorithm can make. This model reveals the minimum number of choices required to correctly distinguish between all possible arrangements of data. It is a fundamental tool for proving that no comparison-based method can ever be faster than a certain mathematical limit for tasks like sorting a list of unique items.

NODIA

170. (B) P-2, Q-1, R-4, S-3
Different ways of exploring networks use different internal storage tools. Layered searches use a first-in, first-out container, while deep-path searches use a last-in, first-out container. Optimization tasks often rely on priority-based structures to always pick the best current option, or specialized tools for managing and merging separate groups of points as they build a minimal connection structure for the entire network.

171. (D) A proposed solution can be checked for correctness in polynomial time.
This category of problems has the defining characteristic that if someone provides a potential answer, it is very easy to verify if that answer is correct. This does not mean the answer itself is easy to find; for many problems in this group, the only known ways to find a solution are slow and involve checking many different possibilities. This distinction is at the heart of major unsolved questions in computing.

172. (B) Fast Fourier Transform (FFT)
Performing calculations on large mathematical expressions can be done much faster by transforming the data into a different numerical form. This transformation allows for much simpler operations on the individual components of the expressions. Once the simplified operations

point. Layered searches are ideal for finding the minimum number of steps in an unweighted path. Ordering tasks based on dependencies is handled by specialized sorting methods. Connecting all points with the lowest total cost is a classic optimization task solved by building a tree one edge at a time.

137. (C) front == -1
In a container that manages elements in a circular fashion, specific indicators are used to track whether any data is currently present. When the pointer responsible for the start of the line is set to a specific non-index value, it signifies that there are no items to process. This simple check is the standard way to confirm that the storage structure is currently empty and ready for new data.

138. (A) Both Statement I and Statement II are true.
Multi-way trees are designed to stay very short by allowing many entries in each node. Because each level can hold many more items than the level above it, the total number of levels increases very slowly as more data is added. This logarithmic growth in height is critical for systems that store massive amounts of information, as it ensures that any piece of data can be reached quickly.

139. (D) Stack
Reversing actions in a text editor requires keeping track of changes in the exact opposite order they were made. The most recent action is always the first one to be undone. A structure that provides access to the most recently added item first is the perfect tool for this functionality. This ensures that the history of changes is managed correctly and that users can easily step back through their work.

140. (A) Both A and R are true and R is the correct explanation of A.
Strategies that only consider the best immediate option can be very fast but sometimes fall into traps. By ignoring the long-term consequences of a choice, these methods may miss the true global optimum that could have been reached by making a less obvious initial decision. While they work perfectly for some problems, they are not a universal solution for all types of complex optimization challenges.

141. (C) O(n)
Organizing a set of items into a specialized hierarchical structure where each parent is larger than its children can be done very efficiently. By starting from the bottom levels and moving upwards, the total amount of work required is proportional to the total number of items. This efficient construction is much faster than repeatedly adding individual items, making it a powerful technique for initializing priority-based data structures.

142. (D) Finding the maximum amount of flow that can be sent from a source to a sink vertex.
This challenge involves determining the greatest possible volume that can move through a network of connections, each with its own capacity limit. It is similar to modeling the movement of water through pipes or data through a communication network. The goal is to coordinate the paths so that the total throughput from the starting point to the destination is as large as possible.

143. (A) At least as hard as the hardest problems in NP.
Problems in this category are defined by their relative difficulty. They have the property that if an efficient way to solve any one of them were found, it would provide an efficient way to solve a whole class of other difficult problems. These challenges are not necessarily solvable or verifiable quickly themselves; they simply represent a benchmark of high computational difficulty in the study of algorithm complexity.

144. (B) Dynamic size; insertions and deletions are efficient.
Connected node structures offer a significant advantage in flexibility. They can grow and shrink as needed without requiring large blocks of contiguous memory. Adding or removing an item at any known point only requires updating a few links, rather than shifting all subsequent items. This makes them much more efficient for applications where the amount of data changes frequently and the order of items is constantly modified.

145. (B) 1, 2, 3, 4
The growth of computational work varies dramatically between different algorithm types. Some processes become impossibly slow even for

Downloaded from www.pdf.tube

standard way to describe these requirements by defining upper and lower limits on their growth. This comprehensive approach allows computer scientists to understand the fundamental efficiency of a solution and to compare different methods fairly on a universal scale.

182. (B) The size is fixed, and it can lead to wasted space or overflow.
Using a simple sequence for storage has limitations when the amount of data changes. If the allocated space is too small, no more items can be added, even if the system has more memory available elsewhere. Conversely, if the space is too large, most of it remains empty and unused. This rigid structure makes it difficult to manage data efficiently in dynamic environments where the number of items varies significantly over time.

183. (A) Both Statement I and Statement II are true.
Decision problems are divided into two fundamental groups based on the type of computer needed to solve them efficiently. One group consists of problems solvable quickly by standard computers. The other group includes problems that can be solved quickly by a theoretical computer capable of exploring many paths simultaneously. This classification is a key concept in the study of computational difficulty and the limits of modern software.

184. (A) O(n^3)
A certain mathematical theorem provides a simple way to find the complexity of a recursive process that divides its work. In this specific case, the number of sub-tasks and how the problem size is reduced lead to a situation where the total work is dominated by the recursive splitting. The result is a complexity that grows very quickly, proportional to the third power of the input size.

185. (C) To find a near-optimal solution for an NP-hard problem in polynomial time.
When a problem is too difficult to solve perfectly in a reasonable amount of time, specialized methods are used to find a result that is very close to the best possible answer. These methods are much faster and are essential for tackling massive real-world problems in logistics, scheduling, and engineering where a perfect solution is less important than finding a very good one quickly and efficiently.

186. (A) n(n-1)/2
In a network where every single pair of points is directly linked, the total number of connections is determined by how many different pairs can be formed from the total population. Mathematically, each point links to every other point, and since each link connects two points, the total count is half the product of the population and the number of other points. This represents the maximum connectivity possible in a simple network.

NODIA

187. (D) Dijkstra's: pick the unvisited vertex with the smallest distance from the source.
Different optimization tasks use different strategies to reach their goals. One common method for connecting all points minimally involves picking the next smallest link that doesn't create a loop. Another approach for finding the best route from a single point involves always picking the closest reachable point. While both use a best-immediate-choice strategy, they are used for solving fundamentally different types of network problems with different final objectives.

188. (C) P-3, Q-4, R-1, S-2
Understanding the efficiency and difficulty of algorithms involves several key concepts. Some methods establish the minimum possible effort for a task, while others categorize the hardest problems whose solutions are easy to check. Analyzing the average cost of a sequence of actions or using chance to improve performance are also important techniques. Together, these tools provide a complete framework for studying how algorithms behave and how they compare.

189. (C) Root, Left, Right
Navigating a hierarchy by visiting the starting point before exploring any of its descendants is a standard way to move through a tree. This method is often used to create a copy of the structure or to display it in a way that shows the parent-child relationships clearly. It ensures that every part of the hierarchy is visited exactly once, starting from the top and moving systematically through each branch.

are complete, the data is transformed back into its original form. This advanced technique is crucial for high-speed processing in signals, images, and other large-scale numerical data applications.

173. (A) Both A and R are true and R is the correct explanation of A.
Storing actual data records only at the bottom level of a hierarchy provides a significant advantage for certain types of queries. By linking all these bottom-level nodes together in a chain, the system can very quickly scan through a whole range of data once it finds the starting point. This makes the structure ideal for database systems where users often need to retrieve many related records in a single operation.

174. (D) Binary Search
Some searching methods achieve their high speed by exploiting the existing order of the data. By always checking the middle item and narrowing the search to one half of the remaining list, they can find a target in a very small number of steps. This efficiency is only possible if the data is already organized in a specific sequence; otherwise, the algorithm would have no way to know which half to ignore.

175. (A) O(n^{2.81})
Standard ways to combine large grids of numbers take a number of steps proportional to the cube of their size. By using a sophisticated recursive approach that requires fewer individual multiplications, this total effort can be reduced. While more complex to implement, this method provides a significant speed advantage for very large grids, which is important for many advanced scientific and engineering simulations in high-performance computing.

176. (D) I, II, III, and IV
Graph exploration methods provide versatile tools for various tasks. Some use a best-at-the-moment approach to find paths, while others explore in layers or follow paths as deep as possible to reveal dependencies and structure. These different strategies can find shortest routes, organize tasks by priority, and connect all points in a network minimally. Together, they are essential for managing and analyzing complex, interconnected data and systems.

177. (A) O(n)
In the least favorable configuration, a search tree can lose its hierarchical structure and stretch out into a single long path. Each node in this path has only one child, effectively making it a simple list. In this state, the depth of the tree is equal to the total number of items stored, which eliminates the speed advantage of searching through a balanced hierarchy and makes all operations much slower.

178. (A) uses a random number at least once during its execution to make a decision.
Some algorithms incorporate an element of chance to improve their average speed or to find solutions for very difficult problems. By making random choices, they can avoid the specific patterns of data that would slow down a deterministic algorithm. This approach can be very powerful, often providing a highly efficient way to handle tasks that would otherwise be computationally expensive in the worst-case scenario.

179. (C) Both the front and the rear
This specialized linear structure offers high flexibility by allowing items to be added or removed from either of its ends. It combines the behaviors of several other basic structures, making it useful for applications where data needs to be managed from both directions, such as undo-redo systems or certain types of work scheduling. This versatility allows it to be used as a general-purpose container for many different kinds of sequential data tasks.

180. (D) Bellman-Ford Algorithm
Networks that include connections with negative values require a robust approach to find the shortest routes. While many fast methods can fail or give wrong answers in this situation, a more deliberate method that repeatedly checks and updates all possible paths can correctly identify the best route. This method is also capable of detecting when a network has loops that would cause path values to decrease indefinitely.

181. (D) I, II, III, IV, and V
Analyzing the performance of algorithms involves measuring both the time they take and the memory they use as the amount of data increases. Mathematical notations provide a

Click Here to Buy book on Amazon

account all the connections and their weights throughout the entire interconnected system.

199. (A) Both Statement I and Statement II are true.
This sorting method is efficient and uses very little extra memory because it reorganizes data directly within its original storage area. However, it does not naturally preserve the relative order of items with identical values, meaning that two equal items might swap positions during the process. This trade-off between memory efficiency and maintaining the original order of equal items is a common consideration when choosing algorithms.

200. (A) When two different keys produce the same hash value.
Ideally, every piece of information should map to its own unique storage slot. However, because there are many more possible pieces of information than there are available slots, it is inevitable that multiple items will occasionally target the same location. This event requires the system to use extra steps to resolve the conflict, which can slow down the process of finding or adding data to the table.

201. (C) n + 1
Every node in a binary hierarchy has two potential links to children. In a structure with many nodes, many of these links will be empty. Mathematically, there is a fixed relationship between the number of nodes and the total number of these empty links. It turns out that there is always exactly one more empty link than there are nodes, regardless of the specific shape or balance of the hierarchy.

202. (A) The Traveling Salesperson Problem (TSP)
Some optimization tasks are so complex that finding a perfect solution is currently impossible for anything but very small examples. In these cases, it is much more practical to use a fast method that provides an answer that is very close to the best possible one. This approach is essential for large-scale real-world problems like planning delivery routes, where a nearly perfect answer found quickly is very valuable.

203. (B) I, II, III, and V only
Strategies that only consider the best immediate option can be very fast but sometimes fall into traps. By ignoring the long-term consequences of a choice, these methods may miss the true global optimum. While they work perfectly for some problems like finding minimum connections or certain pathfinding tasks, they are not a universal solution for all types of complex optimization challenges that exist in computing.

204. (B) A measure of how the running time of an algorithm scales with the size of the input.
This metric describes the fundamental efficiency of a solution by showing how the required time grows as more data is added. It is not about the actual seconds on a specific machine, but rather the mathematical relationship between the amount of work and the data size. This allows for comparing different approaches in a way that is independent of hardware or specific implementation details.

205. (C) Index 3 will have a chain: 3 -> 17 -> 24
In a storage table where conflicts are handled by linking items together, several pieces of information might end up in the same slot. This happens when the mapping function gives them the same initial address. In this specific case, three different numbers all target the same slot. They are simply linked together in a list at that location, which ensures that all of them are still stored and accessible.

206. (A) P-4, Q-1, R-3, S-2
Different mathematical formulas describe the performance of various types of algorithms. Some represent splitting a task in half with linear merging, while others describe linear processes or simple halving of the work. These relationships determine whether an algorithm grows at a log-linear, quadratic, linear, or logarithmic rate, providing a clear way to predict how each one will behave as the amount of data increases significantly.

207. (A) Depth-First Search (DFS)
Exploring a network by following a single path as far as possible before turning back is a standard way to visit every point. This method uses a last-in, first-out behavior to keep track of the path, ensuring that when a dead end is reached, the search reverts to the most recent branching point. It is a powerful tool for understanding the

Downloaded from www.pdf.tube

steps, which perfectly mirrors the backtracking nature of certain traversal algorithms.

281. (A) To connect all vertices in a graph with the minimum total edge weight.
Network design often requires linking all points together in a way that minimizes construction costs. The goal is to create a spanning structure that includes every node while using edges that sum to the smallest possible value. This structure must remain connected and avoid forming any cycles, ensuring efficiency and economy in the overall physical or logical layout.

282. (B) Determining the maximum rate of flow through a network with capacities.
Many systems involve the movement of resources through channels with specific limits. This problem focuses on calculating the highest possible volume that can travel from a start point to an end point without overloading any single path. It is essential for managing traffic, water distribution, and digital data transmission across complex networks with varying throughput capabilities.

283. (B) Finding the most efficient route for a GPS navigation system.
Navigation systems need to identify the quickest or least costly route between two specific locations. While other algorithms connect entire networks, this one focuses on the optimal path between a single pair of points. It considers the weights of individual segments, such as distance or travel time, to provide the user with the most efficient travel plan possible.

284. (A) MST connects all vertices, while Shortest Path connects two specific vertices.
Although both problems involve minimizing edge weights, they serve different fundamental goals. One aims to unify every single node in the graph into a single, cost-efficient structure. The other is concerned only with the best route between a designated starting node and a designated ending node, regardless of whether other nodes in the graph are included.

285. (C) Minimum Spanning Tree
When a planner needs to connect every location in a set using the least amount of material, they are seeking the most efficient spanning structure. This scenario requires that all neighborhoods are reachable but does not focus on a specific path between any two. Minimizing the total length of the pipe for the entire connected network is a classic optimization problem.

286. (D) Problems that can be solved by a deterministic algorithm in polynomial time.
Computational complexity categories group problems by the time needed to solve them. This specific class includes tasks that are considered tractable because they can be resolved efficiently as the input size grows. If a standard computer can find the answer in a reasonable timeframe, the problem is classified here, representing the foundation of efficiently solvable computational tasks.

287. (C) A proposed solution to the problem can be verified in polynomial time.
Some problems might be very difficult to solve, but once a potential answer is provided, checking that answer for correctness is relatively fast. This property of efficient verification is the defining trait of this broad complexity class. It does not guarantee that finding the solution is fast, only that confirming a given solution's validity is computationally manageable.

288. (B) It must be in NP, and it must be NP-hard.
Classification of the most difficult problems in a certain set requires meeting two specific criteria. First, the problem's solution must be verifiable within a reasonable timeframe. Second, it must be at least as difficult as every other problem in that set, meaning any other problem can be converted into it, representing the hardest challenges within the group.

289. (D) It is used to decrease the memory requirements of a problem.
Reducibility is a fundamental concept used to analyze the relationship between different computational problems by transforming one into another. Within the framework of complexity theory, this technique allows researchers to determine relative difficulty and resource needs. Specifically, it can be utilized as a method to decrease the memory requirements needed to solve specific instances of a given problem effectively.

NODIA

217. (C) In-order successor
In a specialized hierarchical structure, unused links are repurposed to point directly to the next item in a specific numerical sequence. This allows the tree to be navigated without needing extra memory to track the path. When reaching the end of a branch, these links provide a shortcut to the next element, making the entire process of visiting every node more efficient and simpler to implement.

218. (A) O(log n)
Adding new data to a self-balancing search tree involves finding the correct spot and then adjusting the structure to maintain equilibrium. Because the tree is kept very short and organized, finding the spot is fast. The adjustments also take a very small amount of work. This ensures that the time required to add an item only grows very slowly as the total number of items in the tree increases.

219. (D) A globally optimal solution can be arrived at by making a sequence of locally optimal choices.
For certain types of optimization tasks, always making the best possible immediate decision is a guaranteed way to reach the best overall result. This property is significant because it allows for very simple and fast algorithms that don't need to consider the future or look back at the past. While not true for all problems, identifying when this property applies is key to efficient algorithm design.

220. (B) Combine
Many powerful problem-solving strategies work by splitting a large task into smaller, independent parts that are solved separately. Once these small solutions are found, they must be brought together to form the final result. this final step is essential for tasks like sorting multiple lists into one or aggregating results from several different sub-processes to reach a single, correct conclusion for the original problem.

221. (C) Finding the Minimum Spanning Tree
Certain optimization tasks can be solved very quickly using efficient, step-by-step strategies. These problems are considered easy because their solutions are found in a reasonable amount of time even for large inputs. In contrast, many other famous challenges currently have no such fast solutions, and are grouped into a class of problems that are significantly harder to solve, even though their results are easy to verify.

222. (D) f(n) is asymptotically less than or equal to g(n).
This mathematical notation provides a way to describe the maximum possible growth rate of an algorithm's resource requirements. It says that for large amounts of data, the actual work will never exceed a certain function multiplied by a constant. This provides a reliable upper bound, allowing developers to know the worst-case performance and ensure that a solution will work within the given constraints.

223. (B) I, II, III, and IV only
Computer science uses various specialized procedures for different tasks. Some focus on identifying text patterns, while others accelerate calculations on mathematical expressions or find properties of numbers like divisors and primality. While many pathfinding methods are efficient, they are categorized differently from these mathematically focused procedures. Together, these tools form a diverse set of solutions for a wide range of computational and analytical problems in modern computing.

224. (B) A queue
Exploring a network layer by layer requires a way to remember which points have been found and are waiting to be visited. A structure that processes items in the exact order they arrive ensures that the search moves outward evenly in all directions. This behavior is essential for finding the minimum number of connections between two points in any unweighted network, providing a structured and reliable search method.

225. (B) The number of zero elements is much greater than the number of non-zero elements.
Grids used in many scientific and engineering tasks often contain a huge amount of empty information. Storing every single cell individually is highly inefficient and wasteful of memory. Specialized storage methods focus only on the significant data points that are not empty. This approach is vital for handling massive datasets where only a tiny fraction of the potential connections or values are actually present.

Click Here to Buy book on Amazon

226. (A) Both Statement I and Statement II are true.
A certain type of multi-way tree is organized so that every single data item is found at the lowest level of the hierarchy. The higher levels are only used as a roadmap to guide the search to the correct location at the bottom. This design is very efficient for database systems, as it allows for both fast searching for individual items and very quick scanning of large ranges of data.

227. (C) Prove that problem A is hard by showing a known hard problem B reduces to it.
Showing that a new task is at least as difficult as a known benchmark of difficulty is a key technique in complexity theory. It involves demonstrating a fast way to translate the difficult task into the new one. If the new task could be solved easily, then the known hard one could also be solved easily, which is a powerful way to categorize the difficulty of different challenges.

228. (D) Bubble Sort
This basic sorting method repeatedly moves through a list, comparing pairs of items that are next to each other. If they are in the wrong order, it swaps them. This process continues until the largest remaining item floats to its correct position at the end. While very simple to understand and implement, it is generally very slow for large amounts of data compared to more advanced methods.

229. (C) Topological sort
Organizing a set of tasks so that every requirement is met before a task begins is a common challenge in planning and scheduling. This results in a sequence where every item that depends on another comes after it in the list. This type of organization is essential for managing complex processes with many dependencies and is only possible in networks that do not have any circular requirements.

230. (B) a b c * +
Hierarchical representations of mathematical expressions can be navigated in several ways. Visiting the components from the bottom up ensures that the values needed for an operation are always ready before the operation is performed. In this specific sequence, the individual values are followed by their corresponding operators, resulting in a notation that is very easy for a computer to evaluate using a simple stack-based system.

231. (D) An optimal solution to the overall problem contains within it optimal solutions to sub-problems.
Efficient problem-solving often relies on a specific property where the best overall result is built from the best results for its smaller parts. This means that if you find the absolute best way to solve a part of the problem, that solution will also be a part of the absolute best solution for the entire thing. Identifying this property is key to using advanced optimization techniques.

232. (A) Contiguous memory locations
This basic storage structure keeps all its items side-by-side in a single block of memory. This organization allows the computer to calculate the exact location of any item instantly using its position in the list. This makes it extremely fast for accessing information, although it also makes it more difficult to grow or shrink the structure once it has been created in memory.

233. (D) n(n-1)/2
In any simple network where links only connect distinct pairs of points without duplicates, there is a maximum possible number of connections. This limit is reached when every point is linked directly to every other point. Mathematically, this is the total number of unique pairs that can be created from the entire population, representing the absolute upper bound on connectivity for that specific number of points.

234. (A) Using a rank-based approach to determine the final position of each element in parallel.
Combining two organized lists can be done much faster by using multiple processors. Instead of one processor doing all the work one step at a time, each unit can simultaneously figure out where its specific items should end up in the final, single list. This coordinated simultaneous work reduces the total time needed to finish the task, taking full advantage of the power of modern multi-processor systems.

235. (D) Decision problems
The classification of the most difficult verifiable problems specifically applies to tasks that have a simple yes or no answer. While many other

NODIA

types of problems exist, these are the foundation for studying computational complexity. Every other verifiable task can be transformed into one of these, providing a clear and standard way to measure and compare the inherent difficulty of many different computational challenges.

236. (B) O(log_t n)
Multi-way trees are designed to stay very short by allowing many branches to originate from each node. Because the capacity of the tree grows exponentially with its depth, the total number of levels required to store a set of items is very small relative to the total count. This ensures that any piece of information can be reached very quickly, even in massive and complex datasets.

237. (C) Every pair of distinct vertices is connected by a unique edge.
A network that has the maximum possible amount of direct connectivity is one where every point is linked to every other point. This ensures that any two points can reach each other in a single step, without needing to pass through any intermediate points. This high level of interconnection is a fundamental property used in theoretical studies of networks, communication patterns, and mathematical graph theory.

238. (D) A highly unbalanced partition and worst-case performance.
The efficiency of a certain recursive sorting method depends on its ability to divide data into roughly equal parts. If the chosen dividing point is consistently near the extreme values of the set, the work is not divided effectively. This results in many more steps being required to finish the task, making the entire process much slower than it would be with better, more balanced divisions.

239. (A) A graph optimization problem
Determining the greatest possible throughput in a network of connections, each with its own limit, is a classic challenge in network analysis and logistics. It involves coordinating the movement of resources so that the total amount reaching the destination is as large as possible. This type of problem is fundamental for managing transportation systems, water distribution, and the efficient flow of data through the internet.

240. (B) Tight asymptotic bound
This specific mathematical notation describes the performance of an algorithm by defining both an upper and a lower limit that it will stay within as the data size grows. It means that the growth rate is accurately captured by a specific function, providing the most precise and complete description of how an algorithm's resource requirements will behave for large amounts of information in real-world use.

241. (B) Deque
This versatile storage structure offers a high degree of flexibility by allowing data to be added or removed from either of its ends. It combines the behaviors of several other simpler structures, making it a powerful tool for various tasks where data management is needed from both directions. This versatility allows it to serve as a general-purpose container for many different kinds of sequential processing.

242. (A) The worst-case number of comparisons.
A diagram representing all possible decisions an algorithm can make reveals its maximum workload. The longest path from the start to an endpoint shows the most steps the algorithm could ever take for any possible input. This is a vital measure for understanding the guaranteed performance of a solution and ensures that it will complete its task within a certain time limit.

243. (B) Binary search
Searching for a value in an organized hierarchy works by repeatedly choosing the correct branch and ignoring the others. This process is very similar to searching in a sorted list by always checking the middle item and narrowing the search to one half. In both cases, each step significantly reduces the remaining area to search, leading to very fast and efficient data retrieval.

244. (D) Its running time is a random variable, but the output is always correct.
Certain problem-solving methods use chance to improve their average speed. While the specific path taken and the time required can vary between different runs, the final result is guaranteed to be accurate every time. This approach is highly effective for tasks where a deterministic method might be very slow in the worst-case, providing a

Downloaded from www.pdf.tube

than specific seconds, developers can compare different algorithms objectively, regardless of the environment in which the computer code is eventually executed or deployed.

263. (C) Big-Theta (Θ)
Different notations describe various limits on the growth of a function. While some provide upper or lower boundaries, this specific notation indicates that the growth rate is bounded from both directions. It implies that the function grows at the same rate as the reference function, providing the most precise description of an algorithm's efficiency for specific input size categories.

264. (D) Recursive algorithms.
Analyzing algorithms that call themselves requires mathematical expressions that define a function in terms of smaller versions of the same problem. These equations describe the total work performed by summing the costs of the sub-problems and the division steps. Solving these relations using various mathematical techniques reveals the overall time complexity of complex, self-referential algorithmic design approaches.

265. (A) It focuses on the growth rate of resource usage relative to input size.
Evaluating how an algorithm performs involves looking at its efficiency as the volume of data increases. The analysis does not seek exact runtimes on specific machines but rather the general trend of resource consumption. This approach allows for scalability comparisons, helping engineers predict how a system will behave when faced with significantly larger datasets in complex production environments.

266. (B) Breaking the problem into smaller sub-problems.
Hierarchical problem-solving strategies begin by decomposing a large, complex task into more manageable parts. These sub-problems must be of the same type as the original to allow for recursive processing. This initial phase is crucial because it sets the foundation for individual solutions that will eventually be merged to form the final result for the original computational problem.

267. (B) To avoid the redundant computation of overlapping sub-problems.
Many complex problems involve sub-tasks that are repeated multiple times throughout the calculation process. Without an efficient strategy, these identical tasks would be recalculated repeatedly, leading to exponential time complexity. This technique optimizes performance by ensuring each sub-problem is solved only once and its result is stored for future reference, significantly improving the overall execution speed.

268. (A) Dynamic Programming has overlapping sub-problems, while Divide and Conquer has independent ones.
The relationship between sub-problems determines which design strategy is most effective. One approach handles tasks that do not share common calculations, allowing them to be solved in isolation. The other is specifically designed for scenarios where the same sub-tasks appear multiple times. Storing results in the latter case prevents unnecessary work, distinguishing it from simpler recursive decomposition methods.

269. (C) Optimal substructure
This property describes a problem where the best possible solution to the whole can be derived from the best possible solutions of its constituent parts. It is a fundamental requirement for applying certain optimization techniques. When this characteristic is present, developers can build up a global solution by making correct, optimal decisions at every lower level of the problem.

270. (D) A typical implementation of Dynamic Programming.
Starting with the smallest possible sub-tasks and using their results to solve increasingly larger instances is a hallmark of efficient optimization. This method avoids the overhead of recursion by populating a table in a structured, sequential manner. It ensures that every necessary piece of information is available before it is needed by a more complex version of the problem.

271. (C) To make a locally optimal choice at each stage of the problem.
This strategy involves selecting the best available option at any given moment without considering the long-term consequences. The goal is to reach

NODIA

Sequential memory structures often require significant shifting of existing elements when items are inserted or removed. Linked structures mitigate this overhead by using pointers to connect separate nodes. Modifying the list only requires updating a few memory addresses rather than rearranging the entire dataset. This characteristic makes them ideal for applications requiring frequent modifications to dynamic data.

255. (C) A Queue, because it processes elements in the order they are received.
Managing tasks based on their arrival sequence requires a First-In, First-Out operational logic. This ensures that the oldest pending request is handled before newer ones, maintaining fairness and order. Systems like print managers rely on this sequential processing to prevent newer jobs from jumping ahead of those already waiting, making it the most appropriate choice for this scenario.

256. (B) To facilitate tree traversal without using a stack or recursion.
Standard tree traversals typically depend on recursive calls or external stacks to keep track of nodes. By replacing null pointers with links to in-order successors or predecessors, these structures enable efficient movement throughout the hierarchy. This modification simplifies the traversal process and reduces memory overhead by eliminating the need for additional data structures during the navigation of the tree.

257. (D) Values in the left subtree are smaller, and values in the right subtree are larger.
Ordered hierarchical structures maintain a specific sorting rule to enable efficient searching. Every node acts as a partition point where all elements in its left branch possess values lower than the parent. Conversely, all elements in the right branch contain values higher than the parent. This consistent organization allows for logarithmic search times across the entire data structure.

258. (A) They guarantee logarithmic time complexity by maintaining height balance.
Standard search trees can become unbalanced, leading to inefficient linear performance during search operations. Self-balancing structures address this by performing rotations to maintain a strict height difference between subtrees. By ensuring the tree remains balanced, these structures guarantee that operations like insertion, deletion, and searching always occur within logarithmic time, even when dealing with the most difficult data inputs.

259. (C) When managing datasets that are too large to fit in main memory.
Massive datasets often require storage on disk, where access times are significantly slower than main memory. Multi-way search structures are designed to minimize these disk read operations by increasing the branching factor of each node. By keeping the overall height of the hierarchy low, these structures allow for efficient data retrieval while processing very large files or databases.

260. (A) stores all data records exclusively in its leaf nodes.
Specialized tree structures optimize range queries by organizing where data is stored. In this specific variation, internal nodes act purely as navigational guides, while the actual data records reside only at the bottom level. These bottom nodes are often linked together, allowing for rapid sequential access across a range of values, which is particularly beneficial for modern database indexing.

261. (D) The amount of memory space it requires to run.
Measuring algorithmic efficiency involves evaluating both time and resource consumption. This specific metric quantifies the maximum amount of working memory required by an algorithm during its execution relative to the input size. It accounts for both static variables and dynamic allocations. Understanding these requirements is essential for developing software that runs effectively on hardware with limited physical memory resources.

262. (B) To provide a machine-independent way of describing an algorithm's growth rate.
Exact execution times fluctuate based on hardware specifications, compiler optimizations, and background processes. Mathematical notations provide a standardized method to describe how resource requirements scale as input sizes increase. By focusing on growth rates rather

Click Here to Buy book on Amazon

much faster and more consistent experience on average for most users.

245. (D) Complexity Theory
The study of how different computational tasks relate to each other and how hard they are to solve relies on several key concepts. One such tool involves showing how one task can be translated into another, allowing researchers to categorize problems into different groups. This provided framework is essential for understanding the fundamental limits of what computers can do and for developing efficient software solutions.

246. (B) Deleting a node whose location is given.
In a chain where each item is linked to both its predecessor and its successor, removing an item is extremely simple if its position is known. You only need to update the links of the two neighboring items to point directly to each other. This is much faster than in other structures where you would have to search from the beginning to find the preceding neighbor.

247. (D) Implementing file systems and databases.
Multi-way trees are the foundation of modern large-scale storage. Their design allows for managing massive amounts of data while keeping the number of steps to find any item very small. This is crucial for systems that store information on slow external disks, as it minimizes the number of times the disk must be accessed, leading to much faster and more reliable overall performance.

248. (A) O(n)
This mathematical relation describes a process that takes a fixed amount of time for each item in a collection. It is characteristic of simple, sequential tasks like visiting every node in a chain once or performing a basic operation on every element in a list. The total work increases directly and steadily as more data is added, resulting in a simple linear relationship.

249. (D) Binary Probing
Mapping data into a table can lead to conflicts where multiple items target the same slot. Several standard methods exist to handle these situations, such as searching for the next empty space or linking conflicting items together. While these established techniques are widely used and effective, some named methods are not part of the standard set of conflict resolution strategies used in computer science.

250. (A) Its running time matches the proven lower bound for the problem for large inputs.
An algorithm is considered as good as it can possibly be when its performance matches the mathematical limit for that specific task. This means that no other possible method, regardless of how clever it is, can be fundamentally faster for large amounts of data. Reaching this benchmark is a major achievement in algorithm design, indicating that the solution is theoretically perfect.

251. (B) It conserves memory by storing only non-zero elements.
Sparse matrices contain many zero elements, making full storage inefficient. Using specialized arrays allows developers to store only non-zero values and their corresponding indices. This approach significantly reduces memory usage compared to storing a complete grid of values. It is a highly effective optimization for large datasets where space conservation is prioritized over raw speed.

252. (A) Stacks
The Last-In, First-Out principle describes a structural behavior where the most recently added item is the first one to be removed. Linear data structures like these are essential for managing function calls, undo operations, and expression evaluations. Unlike queues that prioritize order of arrival, this mechanism focuses on retrieving the latest entry available within the collection.

253. (C) A Priority Queue processes elements based on priority, not just arrival time.
Standard queue structures process items in the order they arrive, following a simple sequential logic. In contrast, specialized variations assign specific values to elements that determine their processing order. Items with higher importance are removed first, regardless of when they entered the system. This allows for more flexible task management in complex environments like operating systems.

254. (D) Their structure allows for element manipulation without shifting other elements.

NODIA

a global solution by consistently choosing the most immediate benefit. While this approach is often fast and simple, it is important to note that it may not always produce the most optimal final outcome possible.

272. (B) It builds a solution incrementally and discards a path once it is deemed non-viable.
Systematic search strategies explore potential solutions step-by-step. When the algorithm realizes that the current path cannot satisfy the necessary constraints, it immediately abandons that branch and returns to a previous state to try a different route. This pruning of the search space allows for the efficient exploration of large sets of possibilities without checking every single invalid combination.

273. (D) To estimate whether a subtree can contain a solution better than the current best.
Optimization searches can be made more efficient by calculating the potential value of a specific search path. If the estimated best outcome from a particular branch is worse than a solution already found elsewhere, that entire branch can be safely ignored. This technique significantly reduces the number of nodes the algorithm must visit, speeding up the discovery of the optimum.

274. (A) Branch and Bound uses a bounding function to prune the search space, which is not a feature of simple backtracking.
Both techniques explore state-space trees to find solutions, but one incorporates an extra layer of logic for optimization. While the simpler method focuses on finding any valid solution by checking constraints, the more advanced version uses mathematical estimates to eliminate paths that cannot possibly improve upon the current best result. This makes it suitable for complex optimization tasks.

275. (C) Backtracking
Finding a solution within a large search space often requires rejecting paths that lead to dead ends. By recognizing constraint violations early, an algorithm can save time by not exploring invalid configurations further. This incremental building and strategic withdrawal is the defining behavior of algorithms designed to navigate complex decision trees where exhaustive searching would be computationally too expensive.

276. (A) A queue
Exploring a network layer by layer requires maintaining a list of nodes in the exact order they were discovered. This ensures that all neighbors at one distance are processed before moving on to nodes further away. A sequential structure that follows the first-in, first-out rule is the most effective way to manage these nodes during the discovery process.

277. (D) It explores as deeply as possible along one path before backtracking.
Graph traversal methods vary in how they move through connections. One specific approach prioritizes depth, following a single branch of nodes until it reaches a dead end or a node with no unvisited neighbors. Only after reaching this point does the algorithm return to the most recent junction to explore other available paths, favoring verticality over breadth.

NODIA

278. (B) Finding the shortest path in an unweighted graph.
Moving outward from a starting point in concentric circles ensures that every node is reached via the minimum number of steps. Because this method explores all immediate neighbors before moving further away, the first time a destination is encountered, it is guaranteed to be through the shortest possible sequence of edges, provided all edge weights are equal.

279. (C) By using recursion or a stack to maintain a LIFO order.
To explore paths deeply, an algorithm must keep track of the path it has taken so it can return to previous nodes. A last-in, first-out mechanism is ideal for this because it allows the algorithm to quickly access the most recent vertex with unexplored branches. This can be handled automatically by the system's call stack or manually via an object.

280. (D) Depth-First Search, as it explores a single path to its conclusion before backtracking.
Solving a maze by picking one direction and following it until you hit a wall aligns with a deep-dive exploration strategy. This method focuses on one possibility at a time rather than trying every nearby path simultaneously. If the path is a dead end, the explorer retraces their

structure and connectivity of complex networks in computing.

208. (B) B+ trees store data pointers only in leaf nodes, while B-trees can store them in internal nodes.
Two common versions of multi-way trees handle data storage in distinct ways. One variant places actual information at every level of the hierarchy, potentially allowing for faster access to some items. The other variant keeps all information at the very bottom level, using the higher levels only as a guide. This second approach is especially useful for scanning through large ranges of data very quickly and efficiently.

209. (A) Both A and R are true and R is the correct explanation of A.
The speed of a certain popular sorting method depends heavily on how well it partitions the data. If the data is already organized, choosing a pivot from the ends can result in very poor partitions where one side is almost empty. This causes the algorithm to take much longer than usual, resulting in a quadratic growth in work that is much slower than its typical fast performance.

210. (A) Representing large graphs with few edges.
Large-scale networks often have very few actual connections compared to the total number of possible links. Storing every single possible link is wasteful. Instead, using a more compact storage method that only records the active connections saves a massive amount of memory. This approach is essential for analyzing massive social networks, web links, or other real-world systems where the number of connections is relatively small.

211. (C) Stack
When a function calls itself, the computer must remember the current state so it can return to it later. This is handled by a container that stores the information for each active call. Because the most recently called function must finish before returning to the previous one, this container must handle items in a last-in, first-out manner, which is the defining behavior of this specific data structure.

212. (D) m
Multi-way trees are defined by a parameter that determines the maximum number of paths that can branch out from a single node. This parameter ensures that the tree stays balanced and that searching through it remains very fast. By allowing many branches at each level, the structure stays very short, making it possible to reach any piece of information in a very small number of steps, even for massive datasets.

213. (A) Both Statement I and Statement II are true.
Some problem-solving methods incorporate an element of chance, meaning they might take different paths or give different results each time they are used. Other methods always follow the exact same steps and will always produce the same result for the same input. Both approaches are valuable in computer science, offering different benefits for performance, reliability, and tackling various types of complex computational challenges.

214. (B) Dynamic Programming
Finding the longest shared sequence between two strings involves solving many smaller versions of the same problem. Because these smaller tasks often overlap, a method that solves each unique sub-task once and saves the answer is very efficient. This approach avoids a massive amount of redundant work, transforming what would otherwise be a very slow process into one that is practical for large strings of information.

215. (D) I, II, III, and IV
Utilizing multiple processing units simultaneously can significantly accelerate computation by dividing a large task into smaller parts. However, this is not always simple; coordinating the different units and merging their results takes time and effort. Not every problem can be divided easily, and the speed improvement is often limited by the parts of the task that must be done one step at a time by a single unit.

216. (A) O(n log n)
While this sorting method can be slow in rare scenarios, it is typically very fast. On average, it partitions data into roughly equal halves at each step, leading to a performance that grows at a log-linear rate. This efficiency makes it a very popular choice for many real-world applications, as it handles most large datasets very effectively and quickly compared to many other common sorting methods.

Downloaded from www.pdf.tube

290. (A) It would mean that P = NP.
The relationship between solvable and verifiable problems is a central mystery in computer science. Because certain difficult problems are linked, finding an efficient way to solve one would immediately provide a way to solve all others in the same group. This would collapse the distinction between these two classes, proving that every problem with a verifiable solution is also efficiently solvable.

291. (C) Number Theoretic Algorithms
Security systems often rely on the mathematical properties of integers to protect data. Tasks such as identifying prime numbers, factoring large values, and performing calculations within modular systems form the backbone of encrypted communication. These specialized mathematical procedures ensure that sensitive information remains secure by creating problems that are easy to set up but extremely difficult for unauthorized parties to solve.

292. (A) Finding occurrences of a pattern within a larger text.
Searching for specific sequences within a vast amount of data is a common requirement in many fields. These algorithms focus on locating every instance where a specific target string appears inside a much longer document or database. This functionality is essential for tools like text editors and search engines, which must process enormous volumes of characters quickly and accurately.

293. (D) The Discrete Fourier Transform (DFT).
Processing signals and multiplying complex mathematical expressions often requires transforming data into a different representation. This specific algorithm provides a highly efficient way to perform this transformation, significantly reducing the number of calculations required compared to traditional methods. Its ability to speed up these operations makes it a foundational tool in modern scientific computing and digital signal processing.

294. (B) Polynomial Arithmetic
Specialized mathematical algorithms are developed to handle operations involving variables raised to various powers. This field includes finding efficient ways to add, divide, and multiply these expressions. By creating optimized procedures for these tasks, developers can solve complex engineering and mathematical problems more quickly, demonstrating the importance of tailoring algorithmic solutions to specific domains of scientific study.

295. (C) String Matching Algorithms
Genomic research involves analyzing massive sequences of genetic information to find specific markers or genes. This task is essentially a search operation where a known pattern is compared against a large dataset. The techniques used to find these occurrences are the same as those used in text processing, highlighting the cross-disciplinary utility of algorithms designed for finding sequences within data.

296. (D) To decrease the total computation time by using multiple processors.
Modern computing hardware often features multiple processing cores that can work on different tasks simultaneously. By designing algorithms that can divide a single problem into independent parts, developers can utilize all available hardware resources. This concurrent execution significantly speeds up the processing of data-intensive tasks like sorting or searching, leading to much shorter overall completion times.

297. (A) When the problem is computationally intractable and an exact solution is not feasible.
Some computational problems are so complex that finding a perfect answer would take years or even centuries. In these cases, it is more practical to use a method that finds a good enough answer quickly. These methods are used when the cost of finding the absolute best solution outweighs the benefits, provided the result is within a guaranteed range.

298. (B) It incorporates an element of randomness into its logic.
Traditional algorithms follow a strictly predictable path, but some modern approaches intentionally use chance to make decisions. By introducing random variables into the computation process, these algorithms can often solve complex problems faster or with less memory on average. While the outcome might vary between executions, the overall efficiency

NODIA

8. What is the primary implication of the equivalence of DFA and NDFA?
(A) Every language accepted by an NDFA is also accepted by a DFA.
(B) DFAs are more powerful than NDFAs in terms of recognition.
(C) NDFAs solve non-computational problems while DFAs cannot.
(D) Converting an NDFA to a DFA results in fewer machine states.

9. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The class of languages accepted by DFAs is the same as the class of languages accepted by NDFAs.
Reason (R) : Every DFA is by definition a special case of an NDFA, and for every NDFA, an equivalent DFA can be constructed using the subset construction method.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

10. Which of the following regular expressions represents the language of all strings over {a, b} that start with 'a' and end with 'b'?
(A) a(a+b)*b
(B) a*b*
(C) (ab)*
(D) a*(a+b)b*

11. A grammar is considered a regular grammar if all its production rules are of the form:
(A) $A \to BC$ or $A \to a$
(B) $A \to \alpha$, where α is any string.
(C) $A \to B$ or $A \to a$
(D) $A \to aB$ or $A \to a$

12. Consider the following statements regarding the properties of regular languages.
(i) The intersection of two regular languages is regular.
(ii) The union of two regular languages is regular.
(iii) The complement of a regular language is regular.
(iv) The concatenation of two regular languages is regular.
Which of the above statement(s) is/are correct?
(A) (i) and (ii) only
(B) (i), (ii), and (iii) only
(C) (ii), (iii), and (iv) only
(D) (i), (ii), (iii), and (iv)

13. The Pumping Lemma for regular languages is primarily used to:
(A) Prove that a given language is regular.
(B) Construct a DFA for a given language.
(C) Prove that a given language is not regular.
(D) Convert an expression into a grammar.

14. Given a language L = {a^n b^n | n ≥ 1}, why is it considered a non-regular language?
(A) It is representable by the expression a*b*.
(B) The language has an infinite string count.
(C) It can be generated by a regular grammar.
(D) It requires counting symbols beyond finite memory.

15. In the context of a compiler, what is the primary function of lexical analysis?
(A) To group characters into meaningful lexemes and produce tokens.
(B) To check the semantic correctness of the source code program.
(C) To generate an optimized machine code for the target processor.
(D) To build a parse tree representing the syntactic structure.

16. Which data structure is central to the operation of a Pushdown Automaton (PDA)?
(A) A queue
(B) A stack
(C) A list
(D) An array

17. Arrange the following language classes of the Chomsky Hierarchy in ascending order of their expressive power (from least powerful to most powerful).
(P) Context-Free Languages
(Q) Recursively Enumerable Languages
(R) Regular Languages
(S) Context-Sensitive Languages
Choose the correct sequence.
(A) R, P, S, Q

Click Here to Buy book on Amazon

CHAPTER 08

Theory of Computation and Compilers

1. Which of the following best defines a formal language in the context of the theory of computation?
(A) A communication system used exclusively by computer scientists.
(B) A subset of strings of symbols derived from a specific alphabet.
(C) A programming language with strictly defined syntax and logic.
(D) A natural language like English that is mathematically modeled.

2. What is the fundamental concept demonstrated by the diagonal argument?
(A) That some problems are computationally unsolvable.
(B) That all infinite sets share the same cardinality.
(C) That the set of real numbers is countably infinite.
(D) That a Turing Machine simulates any other machine.

3. Consider the following statements about non-computational problems. Which statement is correct?
(A) All such problems are solvable with more powerful hardware.
(B) These problems are purely theoretical without practical use.
(C) If no algorithm can solve the problem for all possible inputs.
(D) Any problem stated mathematically is considered computational.

4. Read the following statements regarding set theory and Russell's Paradox.
Statement I : Russell's Paradox arises from considering a set that contains all sets that do not contain themselves.
Statement II : The resolution of Russell's Paradox was a key factor in the development of modern axiomatic set theory.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

5. Which component is NOT part of the 5-tuple formal definition of a Deterministic Finite Automaton (DFA)?
(A) A finite set of states (Q).
(B) A finite set of input symbols (Σ).
(C) A supplementary stack memory unit.
(D) A state transition function (δ).

6. Match List I with List II regarding concepts in the theory of computation.
List I: (Concept)
P. Diagonal Argument
Q. Russell's Paradox
R. Formal Language
S. Non-Computational Problem
List II: (Description)
1. A set of strings over an alphabet.
2. A problem for which no algorithm exists.
3. The set of all sets that do not contain themselves.
4. A proof technique showing some infinite sets are larger than others.
Choose the correct option from those given below:
(A) P-4, Q-3, R-1, S-2
(B) P-3, Q-4, R-2, S-1
(C) P-4, Q-1, R-3, S-2
(D) P-2, Q-3, R-1, S-4

7. A Non-Deterministic Finite Automaton (NDFA) differs from a DFA primarily because its transition function δ maps a state and an input symbol (or ε) to:
(A) Exactly one state.
(B) A symbol on a tape.
(C) The next state in line.
(D) A set of states.

NODIA

frequently surpasses that of purely deterministic methods used today.

299. (B) An Approximation Algorithm.
When a system needs to find a quick solution with a mathematical guarantee of its quality relative to the optimum, it uses specific heuristic techniques. In this scenario, the requirement for the solution to be within 1.5 times the minimum cost indicates that a near-optimal answer is acceptable. This approach allows for efficient operation in complex logistics problems.

300. (A) A lack of guaranteed correctness or variable runtime in exchange for simplicity and efficiency.
Choosing an algorithm that relies on chance involves accepting certain uncertainties. While these methods are often easier to implement and run faster than their deterministic counterparts, they might not yield the same result every time or may take different amounts of time to finish. This balance is often favorable in situations where high speed and simplicity are prioritized.

Downloaded from www.pdf.tube

(B) P, R, S, Q
(C) R, S, P, Q
(D) P, S, R, Q

18. What is the fundamental difference between a deterministic Pushdown Automaton (DPDA) and a non-deterministic Pushdown Automaton (NPDA)?
(A) NPDAs recognize a larger class of languages than DPDAs can.
(B) NPDAs have two stacks, while DPDAs have only one stack.
(C) DPDAs recognize all regular languages, but NPDAs cannot.
(D) NPDAs are equivalent in power to universal Turing Machines.

19. A context-free grammar (CFG) G = (V, T, P, S) is in Chomsky Normal Form (CNF) if every production rule is of the form:
(A) $A \rightarrow aB$ or $A \rightarrow a$
(B) $A \rightarrow aBC$ or $A \rightarrow \varepsilon$
(C) $A \rightarrow BC$ or $A \rightarrow a$
(D) $A \rightarrow Ba$ or $A \rightarrow a$

20. Match List I with List II regarding grammar types and language classes.
List I: (Grammar/Automaton)
P. Finite Automaton
Q. Pushdown Automaton
R. Linear Bounded Automaton
S. Turing Machine
List II: (Language Class Recognized)
1. Context-Sensitive Language
2. Recursively Enumerable Language
3. Regular Language
4. Context-Free Language
Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-3, Q-1, R-4, S-2
(C) P-4, Q-3, R-2, S-1
(D) P-2, Q-4, R-1, S-3

21. A context-free grammar is in Greibach Normal Form (GNF) if all production rules are of the form:
(A) $A \rightarrow BC$, where B and C are non-terminals.
(B) $A \rightarrow a$, where 'a' is a terminal symbol.
(C) $A \rightarrow a\alpha$, where 'a' is a terminal and α is a non-terminal string.
(D) $A \rightarrow \alpha a$, where 'a' is a terminal and α is a non-terminal string.

22. What does it mean for a context-free grammar to be ambiguous?
(A) The grammar generates a language without a PDA.
(B) A string in the language has multiple parse trees.
(C) The grammar contains symbols that are not useful.
(D) The grammar lacks a Chomsky Normal Form variant.

23. A software developer is debugging a parser and needs to visualize how a specific input string id * id + id is derived from the start symbol 'E' of a grammar. Which of the following is the most appropriate representation for this purpose?
(A) A state transition diagram.
(B) A dependency graph.
(C) A control flow graph.
(D) A parse tree.

24. Consider the following statements about the properties of Context-Free Languages (CFLs).
(i) CFLs are closed under union.
(ii) CFLs are closed under intersection.
(iii) CFLs are closed under complementation.
(iv) CFLs are closed under concatenation.
(v) CFLs are closed under Kleene star.
Which of the above statement(s) is/are correct?
(A) (i), (ii), and (iii) only
(B) (ii), (iii), and (v) only
(C) (i), (iv), and (v) only
(D) All are correct.

NODIA

25. The model of a standard Turing Machine consists of an infinite tape, a read/write head, and:
(A) Two stacks and a queue.
(B) A set of registers.
(C) A finite control unit.
(D) A transition function.

26. All of the following are considered variations of the standard Turing Machine that are equivalent in computational power EXCEPT:
(A) Finite Automaton.
(B) Multi-tape machine.
(C) Non-deterministic TM.

(D) Two-way infinite tape.

27. Read the following statements regarding the equivalence of Pushdown Automata (PDAs) and Context-Free Grammars (CFGs).
Statement I : The equivalence between PDAs and CFGs means that for any CFG, a PDA can be constructed that accepts the same language.
Statement II : This equivalence also implies that any language accepted by a PDA can be described by a CFG.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

28. The Church-Turing Thesis proposes that:
(A) Turing machines are the fastest computational model.
(B) Every computable function is computable by a TM.
(C) All problems are solvable by a Universal Machine.
(D) Quantum computers violate computability principles.

29. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A Universal Turing Machine (UTM) can simulate any other Turing Machine.
Reason (R) : A UTM takes as input the description of another Turing Machine (M) and an input string (w), and then simulates the computation of M on w.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

30. A language is said to be Recursive if:
(A) There exists a Turing Machine that loops on reject.
(B) It is generated by a context-free grammar system.
(C) No Turing Machine can be designed for recognition.
(D) A Turing Machine halts on all inputs to accept/reject.

31. Consider the following statements regarding the Chomsky Hierarchy.
(i) Every regular language is also a context-free language.
(ii) Every context-free language is also a context-sensitive language.
(iii) Context-sensitive languages are a proper subset of recursive languages.
(iv) The set of recursive languages is a proper subset of recursively enumerable languages.
Which of the above statement(s) is/are correct?
(A) (i) and (ii) only
(B) (i), (ii) and (iii) only
(C) (i), (ii), (iii), and (iv)
(D) (iii) and (iv) only

32. What is the relationship between Recursive (R) and Recursively Enumerable (RE) languages?
(A) RE is a proper subset of R.
(B) R is a proper subset of RE.
(C) R and RE are disjoint.
(D) R and RE are the same.

33. Which of the following grammars has the most restrictive production rules?
(A) Unrestricted Grammar
(B) Regular Grammar
(C) Context-Sensitive
(D) Context-Free

34. A problem is considered unsolvable or undecidable in the theory of computation if:
(A) It takes exponential time to find any solution.
(B) The problem statement is inherently ambiguous.
(C) No algorithm can solve it for all inputs in finite time.
(D) It requires a non-deterministic machine to solve it.

35. Match List I with List II regarding the decidability status of various problems.
List I: (Problem)
P. Halting Problem for Turing Machines
Q. Membership Problem for Context-Free Languages
R. Emptiness Problem for Finite Automata
S. Post Correspondence Problem
List II: (Decidability Status)

Click Here to Buy book on Amazon

1. Decidable
2. Undecidable
3. Undecidable
4. Decidable

Choose the correct option from those given below:
(A) P-1, Q-2, R-1, S-2
(B) P-2, Q-1, R-4, S-3
(C) P-2, Q-1, R-1, S-2
(D) P-1, Q-2, R-2, S-1

36. An engineer needs to design a Turing Machine to accept the language L = {0^n 1^n | n ≥ 1}. A valid high-level strategy for this TM would be:
(A) Move right and accept if the tape ends immediately.
(B) Compare counts of symbols using a finite state.
(C) Ensure input starts with 0 and ends with 1 only.
(D) Match symbols by replacing them and moving head.

37. The Halting Problem asks whether, for a given arbitrary Turing Machine M and an input w:
(A) The machine M reaches an accepting state.
(B) The machine M produces correct output for w.
(C) The machine M eventually halts its computation.
(D) The machine M uses a finite amount of tape.

38. The Post Correspondence Problem (PCP) is a well-known undecidable problem that involves:
(A) Finding string sequences with matching concatenations.
(B) Determining if two CFGs generate the same language.
(C) Checking if a machine halts on a blank tape input.
(D) Verifying if a regular expression matches an automaton.

39. Which of the following problems is known to be unsolvable (undecidable) for context-free languages?
(A) Checking if the language of a CFG is empty.
(B) Checking if a string is generated by a CFG.
(C) Checking if a given CFG is ambiguous.
(D) Checking if the language of a CFG is finite.

40. Read the following statements regarding computational tractability.
Statement I : Problems that can be solved in polynomial time are considered tractable.
Statement II : Problems that require super-polynomial time (e.g., exponential) are considered intractable.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

41. In the context of parsing, what do associativity and precedence rules determine?
(A) Data types of operands.
(B) Evaluation order.
(C) Memory allocation.
(D) Parsing method.

42. Consider the ambiguous grammar E → E + E | E * E | id. To resolve the ambiguity and enforce standard mathematical precedence, which grammar transformation is appropriate?
(A) E → E + T | T, T → T * F | F, F → id
(B) E → E + E | T, T → T * T | id
(C) E → F + E | F * E | id
(D) E → T + E | T, T → F * T | F, F → id

43. Consider the following statements regarding Top-Down Parsing.
(i) It attempts to construct a parse tree from the root (start symbol) down to the leaves.
(ii) The leftmost derivation of the input string is generated.
(iii) LR parsers are a common example of top-down parsing.
(iv) It may suffer from left-recursion problems in the grammar.
(v) It scans the input from right to left.
Which of the above statement(s) is/are correct?
(A) (i), (ii) and (iv) only
(B) (iii) and (v) only
(C) (i), (iii) and (iv) only
(D) (ii), (iv) and (v) only

44. A key characteristic of a predictive parser, such as an LL(1) parser, is that it:
(A) Requires backtracking to find rules.
(B) Uses lookahead to choose rules.

Downloaded from www.pdf.tube

(C) Builds tree from leaves to root.
(D) Handles any CFG without change.

45. Recursive Descent is a type of top-down parsing that involves:
(A) Mutually recursive procedures for non-terminals.
(B) Using an explicit stack to simulate derivations.
(C) Transforming grammar to an equivalent LR form.
(D) Pre-computing tables before the parsing begins.

46. What does LL(1) stand for in the context of parsing?
(A) Left-to-right scan, Leftmost derivation, 1 symbol lookahead.
(B) Left-to-right scan, Leftmost derivation, 1-ply backtracking.
(C) Left-to-right scan, Rightmost derivation, 1 symbol lookahead.
(D) Lookahead-Left, Leftmost derivation, 1 symbol lookahead.

47. Arrange the following bottom-up parsers in increasing order of their parsing power.
(P) LR(0)
(Q) LALR(1)
(R) LR(1)
(S) SLR(1)
Choose the correct sequence.
(A) P, S, Q, R
(B) P, Q, S, R
(C) S, P, Q, R
(D) P, S, R, Q

48. A major advantage of an LALR(1) parser compared to an LR(1) parser is that it:
(A) Uses fewer states in its table.
(B) Parses wider grammar ranges.
(C) Operates without using a stack.
(D) Is a top-down parsing method.

49. Bottom-up parsing is also known as:
(A) Recursive descent parsing.
(B) Shift-reduce parsing.
(C) Predictive parsing.
(D) Backtracking parsing.

50. Match List I with List II regarding parsing techniques and their descriptions.
List I: (Parsing Term)
P. Top-Down Parsing
Q. Bottom-Up Parsing
R. LL(1)
S. LALR(1)
List II: (Description)
1. A predictive parser with one lookahead symbol.
2. Starts from the input string and tries to reduce it to the start symbol.
3. Merges states of an LR(1) parser to reduce table size.
4. Starts from the start symbol and tries to derive the input string.
Choose the correct option from those given below:
(A) P-2, Q-4, R-3, S-1
(B) P-4, Q-2, R-1, S-3
(C) P-4, Q-2, R-3, S-1
(D) P-2, Q-4, R-1, S-3

51. What is the main purpose of semantic analysis in a compiler?
(A) To convert code into tokens.
(B) To check the syntactic rules.
(C) To verify static logic and types.
(D) To optimize intermediate code.

52. In an attribute grammar, what is a synthesized attribute?
(A) An attribute whose value is computed from its children.
(B) An attribute whose value is computed from its parent.
(C) An attribute that is constant across the entire tree.
(D) An attribute passed from a sibling node in the tree.

NODIA

53. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The Halting Problem for Turing Machines is undecidable.
Reason (R) : If it were decidable, one could construct a new machine that contradicts its own behavior by doing the opposite of what the decider predicts, leading to a paradox.
(A) Both A and R are true and R is the correct

explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

54. An inherited attribute is one whose value at a parse tree node is computed based on the attributes of its:
(A) Parent and/or siblings.
(B) Only its child nodes.
(C) Children and parent.
(D) Leftmost child only.

55. A Syntax Directed Definition (SDD) that uses only synthesized attributes is known as:
(A) An L-attributed definition.
(B) An S-attributed definition.
(C) An ambiguous definition.
(D) A circular definition.

56. Consider the following statements regarding Syntax Directed Definitions (SDDs).
(i) SDDs are context-free grammars with attributes and semantic rules.
(ii) They provide a formal way to specify the semantics of a programming language.
(iii) They can be used to describe both syntax analysis and semantic analysis.
(iv) All SDDs can be evaluated during a single pass of a bottom-up parser.
Which of the above statement(s) is/are correct?
(A) (i) and (ii) only
(B) (i), (ii) and (iii) only
(C) (iii) and (iv) only
(D) (i), (ii), (iii), and (iv)

57. A dependency graph in semantic analysis is used to:
(A) Show the source flow.
(B) Map function calls.
(C) Map stack allocation.
(D) Order attribute evaluation.

58. What does an L-attributed definition allow for attribute computation at a node?
(A) Synthesized attributes only.
(B) Parent and child attributes.
(C) Parent or left-sibling data.
(D) Only parent or left siblings.

59. The primary role of type-checking during the semantic analysis phase is to:
(A) Reduce total memory usage.
(B) Verify operand compatibility.
(C) Generate intermediate code.
(D) Allocate variable storage.

60. Which of the following is NOT typically part of the run-time storage organization for a block-structured language?
(A) Code segment (text)
(B) Stack
(C) Heap
(D) Parse tree

61. What does an activation tree represent?
(A) Dynamic function call sequences.
(B) Program syntactic structures.
(C) Data flow between modules.
(D) Attribute dependencies.

62. An activation record (or stack frame) is used to store information for a single invocation of a procedure. Which of the following is typically NOT found in an activation record?
(A) Local variable values.
(B) The return address.
(C) Global symbol table.
(D) Passed parameters.

63. The use of a stack for allocating activation records is suitable for programming languages that feature:
(A) Recursion.
(B) Goto jumps.
(C) Parallelism.
(D) Self-modifying code.

64. In which parameter passing mechanism is the address of the argument passed to the procedure, allowing the procedure to modify the original argument?
(A) Pass-by-reference.
(B) Pass-by-value.
(C) Pass-by-result.
(D) Pass-by-name.

65. The primary function of a symbol table in a compiler is to:
(A) Track identifiers and scopes.
(B) Record activation trees.

Click Here to Buy book on Amazon

(C) Manage the memory heap.
(D) Store intermediate code.

66. Which of the following is a common form of machine-independent intermediate representation?
(A) Three-address code.
(B) Target assembly code.
(C) Executable machine code.
(D) Stream of lexer tokens.

67. The translation of a declaration like `int x, y;` by the intermediate code generator primarily involves:
(A) Making arithmetic instructions.
(B) Emitting jump code logic.
(C) Entering data in symbol table.
(D) Optimizing register usage.

68. Consider the assignment statement `a = b + c * d;`. A possible three-address code representation for this would be:
(A) t1 = c * d; t2 = b + t1; a = t2;
(B) t1 = b + c; t2 = t1 * d; a = t2;
(C) a = b + c * d;
(D) mult c, d, t1; add b, t1, t2; move t2, a;

69. How are short-circuiting boolean expressions like `(A and B)` typically translated into intermediate code with control flow?
(A) Compute A and B always.
(B) Jump if A is false.
(C) Jump if A is true.
(D) Ignore A and evaluate B.

70. The translation of a procedure call typically involves generating code for all of the following actions EXCEPT:
(A) Pushing the arguments.
(B) Saving return address.
(C) Transferring control.
(D) Re-parsing the body.

71. A control-flow graph (CFG) in the context of code optimization is a directed graph where:
(A) Nodes are blocks and edges are flow.
(B) Nodes are calls and edges are links.
(C) Nodes are data and edges are deps.
(D) Nodes are trees and edges are links.

72. Data-flow analysis is a technique used to:
(A) Collect possible program values.
(B) Check syntactic correctness.
(C) Manage the runtime stack.
(D) Check type compatibility.

73. Which type of optimization is typically performed on a single basic block of code?
(A) Local optimization.
(B) Global optimization.
(C) Loop optimization.
(D) Interprocedural.

74. Loop unrolling is a loop optimization technique that:
(A) Deletes the entire loop block.
(B) Replicates the loop body.
(C) Moves invariants outside.
(D) Merges nested loop code.

75. Peephole optimization is a technique that improves code by:
(A) Replacing instruction sequences.
(B) Reordering large code blocks.
(C) Analyzing the global graph.
(D) Optimizing between calls.

76. Instruction scheduling is a code optimization phase that aims to:
(A) Reduce total instruction count.
(B) Minimize pipeline stall issues.
(C) Allocate program variables.
(D) Remove loop computations.

77. Which of the following identifies an incorrect pair of a concept and its primary use?
(A) Pumping Lemma - Non-regularity.
(B) CNF - Algorithmic parsing.
(C) Lexical Analysis - Semantic logic.
(D) Symbol Table - Identifier data.

NODIA

78. Read the following statements regarding Deterministic and Non-deterministic Finite Automata.
Statement I : A deterministic finite automaton (DFA) can have multiple transitions from a single state on the same input symbol.
Statement II : A non-deterministic finite automaton (NDFA) can have epsilon (ε) transitions.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.

(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

79. Consider a language L consisting of all strings over {0, 1} with an equal number of 0s and 1s. This language is:
(A) Regular
(B) Context-Free only
(C) Context-Sensitive only
(D) Recursively Enumerable

80. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : S-attributed definitions can be evaluated by a bottom-up parser during parsing.
Reason (R) : In an S-attributed definition, all attributes are synthesized, meaning their values can be computed from the attributes of the children after the children's subtrees have been fully processed.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

81. A grammar G has the productions S → aS | bS | ε. What language does this grammar generate?
(A) Strings ending with a or b.
(B) All strings over {a, b}.
(C) The empty language set.
(D) Only the empty string.

82. The process of converting an NDFA with N states to an equivalent DFA can, in the worst case, result in a DFA with how many states?
(A) 2^N states.
(B) N states.
(C) N^2 states.
(D) log(N) states.

83. Consider the following statements regarding Turing Machines.
(i) A Turing Machine can solve the Halting Problem.
(ii) A Non-deterministic Turing Machine is more powerful than a Deterministic Turing Machine.
(iii) The Church-Turing thesis states that any solvable problem has a Turing Machine solution.
(iv) A Universal Turing Machine can compute anything that any other Turing Machine can compute.
Which of the above statement(s) is/are correct?
(A) (i) and (ii) only
(B) (iii) and (iv) only
(C) (i), (ii) and (iii) only
(D) (i), (iii) and (iv) only

84. In LR parsing, what does a shift action signify?
(A) Successful string recognition.
(B) Pushing symbols onto stack.
(C) Handle reduction process.
(D) Detection of syntax errors.

85. In a compiler, if the expression is `x + y * z`, which operation is typically performed first by the intermediate code generator, assuming standard operator precedence?
(A) Reading the x value.
(B) Addition of x and y.
(C) Multiplication of y, z.
(D) Result assignment.

86. Match List I with List II regarding code optimization techniques.
List I: (Optimization Technique)
P. Loop Invariant Code Motion
Q. Constant Folding
R. Strength Reduction
S. Dead Code Elimination
List II: (Description)
1. Replacing a costly operation (e.g., multiplication) with a cheaper one (e.g., addition).
2. Removing code that is never executed or whose result is never used.
3. Moving a calculation that produces the same result in every loop iteration to before the loop.
4. Evaluating expressions with constant operands at compile time.
Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-2, S-1
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-4, R-1, S-3

Downloaded from www.pdf.tube

program into machine code before execution begins.
(ii) An interpreter executes source program instructions one by one.
(iii) Compiled programs generally run faster than interpreted programs.
(iv) Interpreters offer better portability and debugging capabilities.
(v) Lexical and syntax analysis are performed in both compilers and interpreters.

Which of the above statement(s) is/are correct?
(A) (i), (ii), and (iii) only
(B) (i), iii), and (v) only
(C) (ii), (iv), and (v) only
(D) All are correct.

134. Which of the following is the most powerful automaton?
(A) Finite Automaton
(B) Pushdown Automaton
(C) Turing Machine
(D) Deterministic PDA

135. In a syntax-directed translation scheme, if the semantic rules are placed only at the end of the right-hand side of productions, the scheme is known as:
(A) Postfix translation.
(B) Inherited scheme.
(C) L-attributed scheme.
(D) Circular definition.

136. Which phase of the compiler is primarily responsible for detecting an error like type mismatch in assignment?
(A) Lexical Analysis
(B) Syntax Analysis
(C) Semantic Analysis
(D) Code Generation

137. A three-address instruction has at most:
(A) One op, three addresses.
(B) Three ops, one address.
(C) Two ops, two addresses.
(D) Three ops, three addresses.

138. The pumping lemma is a property that holds for a certain class of languages. If a language fails to satisfy the pumping lemma, what can be concluded?
(A) Not in that class.
(B) Belongs to class.
(C) Test is useless.
(D) It is undecidable.

139. Match List I with List II regarding compiler terms and their associated concepts.
List I: (Term)
P. Backpatching
Q. Left Factoring
R. Peephole Optimization
S. Subset Construction
List II: (Associated Concept)
1. Grammar transformation for predictive parsing.
2. Converting an NDFA to a DFA.
3. Generating code for boolean expressions and flow-of-control.
4. Local code optimization technique.

Choose the correct option from those given below:
(A) P-1, Q-3, R-4, S-2
(B) P-3, Q-1, R-2, S-4
(C) P-3, Q-1, R-4, S-2
(D) P-4, Q-2, R-1, S-3

NODIA

140. A language is accepted by a non-deterministic pushdown automaton (NPDA) but not by any deterministic pushdown automaton (DPDA). What kind of language is it likely to be?
(A) Regular language.
(B) Ambiguous CFL.
(C) Context-sensitive.
(D) Recursive language.

141. The run-time system component responsible for managing memory for objects whose lifetimes are not tied to procedure activations (e.g., via `malloc` or `new`) is the:
(A) Static area.
(B) Heap manager.
(C) Stack manager.
(D) Code segment.

142. A shift-reduce parser has a conflict where, for a given state and lookahead symbol, both a shift action and a reduce action are possible. This is called a:
(A) Shift-reduce conflict.
(B) Shift-shift conflict.
(C) Reduce-reduce conflict.
(D) Accept-reject conflict.

99. A parse tree is a graphical representation of a derivation. What does a leaf node in a parse tree represent?
(A) Non-terminal.
(B) Terminal or ε.
(C) Start symbol.
(D) Production rule.

100. Consider the phases of a compiler listed below.
(P) Semantic Analysis
(Q) Code Generation
(R) Lexical Analysis
(S) Syntax Analysis
What is the correct order of these phases in a typical compiler?
(A) R, S, P, Q
(B) S, R, P, Q
(C) R, P, S, Q
(D) S, P, R, Q

101. An intractable problem is one that is:
(A) Unsolvable by computer.
(B) Not solvable in polynomial.
(C) Only for non-deterministic.
(D) Solvable in linear time.

102. What is a handle in the context of bottom-up parsing?
(A) Substring ready for reduction.
(B) First input character.
(C) Initial reduction rule.
(D) Final input character.

103. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Left recursion in a grammar is problematic for top-down parsers like recursive descent.
Reason (R) : A left-recursive production rule causes the parser to enter an infinite loop of recursive calls without consuming any input.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

104. The storage class for variables that are allocated on the heap at run-time is typically referred to as:
(A) Static allocation.
(B) Automatic allocation.
(C) Dynamic allocation.
(D) Register allocation.

105. When is an S-attributed definition used most effectively?
(A) With top-down parsing.
(B) With bottom-up parsing.
(C) With inherited attributes.
(D) With ambiguous grammar.

106. The language L = {ww | w is in {a,b}*} is not context-free. This can be proven using:
(A) Pumping lemma for CFLs.
(B) Pumping lemma for regular.
(C) Non-acceptance by DFA.
(D) Russell's Paradox logic.

107. Consider the following statements about parsing techniques.
(i) LL parsers are top-down parsers.
(ii) LR parsers are bottom-up parsers.
(iii) Recursive descent parsing is a type of LR parser.
(iv) LR parsers are generally more powerful than LL parsers.
(v) LALR parsers have fewer states than LR(1) parsers for the same grammar.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iii), and (iv) only
(B) (i), (ii), (iv), and (v) only
(C) (ii), (iii), and (v) only
(D) All are correct.

108. Which problem is an example of an unsolvable problem for Turing machines, but a solvable (decidable) problem for finite automata?
(A) The equivalence problem.
(B) The membership problem.
(C) The emptiness problem.
(D) The halting problem.

109. What is a basic block in the context of intermediate code?
(A) Continuous code without jumps.
(B) Single 3-address line.
(C) Whole function body.

Click Here to Buy book on Amazon

(D) Code run at least once.

110. The regular expression `(a|b)*aa(a|b)*` represents the language of all strings over {a, b} that:
(A) Start with 'aa'.
(B) Contain 'aa'.
(C) End with 'aa'.
(D) Have two 'a's.

111. If a language L is recursive, then its complement, L', is:
(A) Also recursive.
(B) Only enumerable.
(C) Not enumerable.
(D) Context-free.

112. A dependency graph with a cycle indicates that:
(A) Grammar is ambiguous.
(B) Attributes are uncomputable.
(C) Order is well-defined.
(D) Infinite code loops.

113. What is the primary advantage of using Greibach Normal Form (GNF)?
(A) It eliminates ambiguity.
(B) Most compact representation.
(C) Step consumes input.
(D) Simplifies LR parsers.

114. Match List I with List II regarding compiler phases and their primary outputs.
List I: (Compiler Phase)
P. Lexical Analyzer
Q. Syntax Analyzer
S. Semantic Analyzer
T. Intermediate Code Generator
List II: (Primary Output)
1. Parse Tree
2. Annotated Parse Tree
3. Stream of Tokens
4. Three-Address Code
Choose the correct option from those given below:
(A) P-3, Q-1, S-2, T-4
(B) P-1, Q-3, S-2, T-4
(C) P-3, Q-2, S-1, T-4
(D) P-1, Q-2, S-3, T-4

115. A grammar is not suitable for LL(1) parsing if it is:
(A) Right recursive.
(B) Left recursive.
(C) Unambiguous.
(D) In CNF form.

116. Consider the process of compiling a program for a machine with a limited number of registers. The phase most directly responsible for assigning variables to these registers is:
(A) Semantic Analysis
(B) Intermediate Code
(C) Code Optimization
(D) Syntax Analysis

117. Read the following statements regarding context-free languages.
Statement I : The language {a^n b^n c^n | n ≥ 1} is a context-free language.
Statement II : The language {a^n b^n | n ≥ 1} is a context-free language.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

118. Which of the following is NOT a form of intermediate representation?
(A) Syntax Tree
(B) Postfix Notation
(C) Three-Address Code
(D) Symbol Table

119. A key difference between global optimization and local optimization is that:
(A) Globals cover procedures.
(B) Locals are machine-tied.
(C) Globals are by humans.
(D) Locals are more effective.

120. Which of the following problems about Turing Machines is decidable?
(A) Halt on input 'w'.
(B) Halt on blank tape.
(C) Halt within 100 steps.
(D) Accept any string.

121. What does the activation tree for a program execution represent?
(A) Static nested structure.
(B) Dynamic call nesting.
(C) Inheritance hierarchy.
(D) Variable data flow.

122. In the context of the Chomsky hierarchy, unrestricted grammars generate which class of languages?
(A) Regular languages
(B) Context-free
(C) Context-sensitive
(D) Enumerable

123. A grammar has the productions: S -> A, A -> A+A | id. A parser for this grammar is processing the input 'id+id+id'. The substring 'A+A' on top of the parse stack is a candidate for reduction. This substring is known as a:
(A) Token
(B) Handle
(C) Lookahead
(D) Non-terminal

124. Consider the following steps in the lifecycle of a variable in a procedure.
(P) Deallocation of memory
(Q) Use of the variable
(R) Declaration of the variable
(S) Allocation of memory
What is a typical chronological order for these steps?
(A) R, S, Q, P
(B) S, R, Q, P
(C) R, Q, S, P
(D) P, S, R, Q

125. The task of resolving an operator like '+' to mean integer addition in one context and floating-point addition in another is handled during:
(A) Lexical Analysis
(B) Syntax Analysis
(C) Semantic Analysis
(D) Code Optimization

126. A language L is recursively enumerable but not recursive. What can be concluded about the complement of L?
(A) It is recursive.
(B) It is enumerable.
(C) It is not enumerable.
(D) It is context-sensitive.

127. Which of the following correctly pairs a parsing technique with its category?
(A) Recursive - Bottom-up
(B) Shift-Reduce - Top-down
(C) LL(1) - Bottom-up
(D) LR(1) - Bottom-up

128. An optimization that replaces the expression `2 * x` with `x + x` is an example of:
(A) Constant folding
(B) Dead code removal
(C) Strength reduction
(D) Loop unrolling

129. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The intersection of two context-free languages is not necessarily context-free.
Reason (R) : The language L = {a^n b^n c^n | n ≥ 1}, which is not context-free, can be expressed as the intersection of two context-free languages, L1 = {a^n b^n c^m | n, m ≥ 1} and L2 = {a^m b^n c^n | m, n ≥ 1}.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

130. The access link (or static link) in an activation record is used to:
(A) Reach non-local data.
(B) Point to the caller.
(C) Hold return address.
(D) Manage heap data.

131. Which of the following statements about formal languages is incorrect?
(A) Regulars are context-free.
(B) CFLs close on intersection.
(C) Regulars close on complement.
(D) CFL unions are context-free.

132. A grammar that produces more than one parse tree for some sentence is called:
(A) Unambiguous
(B) Ambiguous
(C) Context-sensitive
(D) Regular

133. Consider the following statements about compilers and interpreters.
(i) A compiler translates the entire source

NODIA

(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

154. What is the role of a display in some run-time systems?
(A) Accessing non-local data.
(B) Showing user errors.
(C) Stack visualization.
(D) GUI window manager.

155. Which of the following is a key characteristic of an unrestricted grammar?
(A) Rules have no limits.
(B) Only context-sensitive.
(C) Rules are A -> a.
(D) Power of a PDA.

156. An engineer wants to optimize a loop by moving a calculation `x = y * z` out of the loop. This is a valid optimization only if:
(A) Values stay constant.
(B) High iteration count.
(C) Integer data types.
(D) x unused after loop.

157. Which of the following statements about formal language theory is true?
(A) RE complement is RE.
(B) Halting is an RE task.
(C) RE and co-RE is Rec.
(D) Decidable is in P.

158. Which of the following is NOT a task of the lexical analyzer?
(A) Checking parenthesis.
(B) Removing whitespace.
(C) Finding identifiers.
(D) Creating token streams.

159. Consider a syntax-directed definition for a simple desk calculator. The attribute associated with a non-terminal 'E' representing an expression would most likely be:
(A) Numerical value.
(B) Data type tag.
(C) Operator count.
(D) Line position.

160. What is the fundamental limitation of a finite automaton that prevents it from recognizing a language like L = {a^n b^n | n ≥ 0}?
(A) No non-determinism.
(B) Limited state count.
(C) Read-only tape.
(D) No accepting states.

161. In the context of computational complexity, a problem that is known to be solvable but for which no polynomial-time algorithm has been found is classified as:
(A) Intractable.
(B) Unsolvable.
(C) Tractable.
(D) Non-computational.

162. Match List I with List II regarding grammar types and their corresponding automata.
List I: (Grammar Type)
P. Type 0 (Unrestricted)
Q. Type 1 (Context-Sensitive)
R. Type 2 (Context-Free)
S. Type 3 (Regular)
List II: (Corresponding Automaton)
1. Pushdown Automaton
2. Finite Automaton
3. Turing Machine
4. Linear Bounded Automaton
Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-1, S-2
(C) P-3, Q-4, R-2, S-1
(D) P-2, Q-1, R-4, S-3

163. The process of modifying a grammar by replacing a production $A \rightarrow A\alpha \mid \beta$ with $A \rightarrow \beta A'$ and $A' \rightarrow \alpha A' \mid \varepsilon$ is called:
(A) Elimination of left recursion.
(B) Left factoring.
(C) Production substitution.
(D) Chomsky normalization.

164. What is the primary purpose of a return address in an activation record?
(A) Resuming caller code.
(B) Storing call result.
(C) Target machine code.
(D) Heap start address.

Downloaded from www.pdf.tube

143. Consider the following statements about code optimization.
(i) Optimization can sometimes increase the size of the code.
(ii) Peephole optimization is a type of global optimization.
(iii) Loop-invariant code motion can improve program efficiency significantly.
(iv) Optimization must not change the semantic meaning of the program.
Which of the above statement(s) is/are correct?
(A) (i), (iii) and (iv) only
(B) (ii) and (iv) only
(C) (i), (ii) and (iii) only
(D) All are correct.

144. What is the key idea behind a Universal Turing Machine (UTM)?
(A) It simulates any TM.
(B) Solves undecidables.
(C) Faster than standard.
(D) Physical hardware.

145. Call-by-name is a parameter passing mechanism where:
(A) Text is substituted.
(B) Copy is passed.
(C) Address is passed.
(D) Evaluated once.

146. Read the following statements regarding LALR(1) and LR(1) parsers.
Statement I : An LR(1) parser is more powerful than an LALR(1) parser.
Statement II : An LALR(1) parser may have reduce-reduce conflicts that are not present in the corresponding LR(1) parser.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

147. A DAG (Directed Acyclic Graph) is often used in compilers as a more compact representation of expressions than a syntax tree because:
(A) It handles common parts.
(B) Shows ambiguities.
(C) Allows cyclic loops.
(D) Easier code output.

148. Which of the following is NOT a property under which the class of regular languages is closed?
(A) Union
(B) Subset
(C) Intersection
(D) Complementation

149. The problem of determining whether two arbitrary context-free grammars generate the same language is:
(A) Decidable
(B) Undecidable
(C) Tractable
(D) Polynomial

150. Arrange the following compiler-related activities in the order they would typically occur.
(P) Register Allocation
(Q) Type Checking
(R) Tokenization
(S) Parse Tree Construction
Choose the correct sequence.
(A) Q, R, S, P
(B) R, S, Q, P
(C) R, Q, S, P
(D) S, R, Q, P

151. What does the Church-Turing Thesis postulate?
(A) TMs compute everything.
(B) Universe is a machine.
(C) All problems are easy.
(D) Non-det is powerful.

152. In a bottom-up LR parser, the 'reduce' action involves:
(A) Handle swap for non-terminal.
(B) Pushing new symbol.
(C) Shifting input token.
(D) Reporting an error.

NODIA

153. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Regular languages are a proper subset of context-free languages.
Reason (R) : Every DFA can be simulated by a PDA that ignores its stack, and there exist context-free languages, like {a^n b^n | n ≥ 1}, that are not regular.
(A) Both A and R are true and R is the correct explanation of A.

(C) Syntax Analysis.
(D) Semantic Analysis.

176. Which of the following is an unsolvable problem related to context-free languages?
(A) Ambiguity check.
(B) Membership check.
(C) Emptiness check.
(D) Finiteness check.

177. A grammar `G` is defined by the productions `S → aSb | ε`. What is the language `L(G)`?
(A) {a^n b^n | n ≥ 0}
(B) {a^n b^m | n, m ≥ 0}
(C) {(ab)^n | n ≥ 0}
(D) {a, b}*

178. The control link (or dynamic link) in an activation record typically points to:
(A) Record of the caller.
(B) Enclosing scope link.
(C) Global data region.
(D) Memory heap region.

179. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The class of context-free languages is closed under union.
Reason (R) : If L1 and L2 are context-free languages with start symbols S1 and S2, a new grammar can be constructed with a new start symbol S and productions S -> S1 | S2, which generates the union L1 ∪ L2.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

180. The primary goal of instruction scheduling is to:
(A) Minimize stall issues.
(B) Lower instruction count.
(C) Map variables to reg.
(D) Remove redundant code.

181. A student is tasked with creating a parser for a grammar. They choose an LR(1) parser. What does the 1 in LR(1) signify?
(A) Single symbol lookahead.
(B) Only one grammar type.
(C) Single backtrack count.
(D) Single tree level step.

182. Which of the following correctly describes a non-computational problem?
(A) No algorithm can exist.
(B) Long computation time.
(C) Theoretically solvable.
(D) Only quantum solvable.

183. In the context of semantic analysis, which of the following is an example of an inherited attribute?
(A) Downward type passing.
(B) Result of expression.
(C) Identifier character.
(D) Total node count sum.

184. The regular expression for a language consisting of strings with an even number of 'a's over the alphabet {a, b} is:
(A) (b*ab*a)*b*
(B) (a*ba*b)*
(C) (ab)*
(D) a(ba)*b

185. What is the Post Correspondence Problem (PCP)?
(A) Undecidable string match.
(B) Shortest path graph.
(C) Decidable email logic.
(D) Optimization problem.

186. Which of the following is true about LALR(1) parsers?
(A) Merged core LR(1) states.
(B) More powerful than LR(1).
(C) State count of SLR(1).
(D) Never have SR conflicts.

187. Consider the following statements regarding formal grammars.
(i) A context-free grammar is more powerful than a regular grammar.
(ii) The production A -> ε is allowed in any context-sensitive grammar.
(iii) Unrestricted grammars generate recursively enumerable languages.
(iv) A Chomsky Normal Form grammar has rules of the form A -> aB or A -> a.

Click Here to Buy book on Amazon

165. Read the following statements regarding Peephole optimization.
Statement I : Peephole optimization is applied to the final machine code.
Statement II : Peephole optimization is a form of local optimization that works on a small, fixed-size window of instructions.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

166. A reduce-reduce conflict occurs in an LR parser when:
(A) Choice between reductions.
(B) Shift versus reduce.
(C) Rule not in grammar.
(D) Invalid input token.

167. Which of the following statements correctly describes the relationship between L-attributed and S-attributed definitions?
(A) All S are also L.
(B) All L are also S.
(C) Classes are disjoint.
(D) Classes are identical.

168. The Halting Problem is undecidable, which implies that:
(A) No universal decider exists.
(B) TMs cannot be built.
(C) Programs always halt.
(D) NP class is unsolvable.

169. Consider the following statements regarding compiler construction.
(i) A symbol table is typically a data structure that persists through all phases of compilation.
(ii) The output of the syntax analysis phase is a parse tree.
(iii) Type checking is a responsibility of the syntax analyzer.
(iv) Peephole optimization examines a small window of instructions to find improvements.
(v) LALR parsers are less powerful than SLR parsers.
Which of the above statement(s) is/are correct?
(A) (i), (ii), and (iv) only
(B) (i), [iii), and (v) only
(C) (ii), (iv), and (v) only
(D) All are correct.

170. If a language is regular, it can be described by all of the following EXCEPT:
(A) Regular expression.
(B) A deterministic DFA.
(C) A non-deterministic.
(D) Complex stack PDA.

NODIA

171. In data-flow analysis, the concept of reaching definitions is used to determine:
(A) Unchanged assignments.
(B) Live program points.
(C) Common expressions.
(D) Control flow logic.

172. The Chomsky Hierarchy classifies formal languages based on the restrictiveness of the grammar rules required to generate them. Which is the correct hierarchy from most restrictive (least powerful) to least restrictive (most powerful)?
(A) Regular, Context-Free, Context-Sensitive, Unrestricted
(B) Context-Free, Regular, Context-Sensitive, Unrestricted
(C) Regular, Context-Sensitive, Context-Free, Unrestricted
(D) Unrestricted, Context-Sensitive, Context-Free, Regular

173. When translating control flow statements like an `if-then-else` statement, an intermediate code generator will typically make use of:
(A) Labels and jump code.
(B) Only basic arithmetic.
(C) Stack push and pop.
(D) Procedure call rules.

174. A grammar is left-factored to:
(A) Aid predictive choice.
(B) Remove left recursion.
(C) Remove ambiguity logic.
(D) Form Chomsky grammar.

175. Which of the following phases of a compiler is LEAST likely to be found in a simple interpreter?
(A) Optimization phase.
(B) Lexical Analyzer.

87. Which of the following is the best description of a context-sensitive grammar?
(A) A → aB or A → a rules.
(B) Single non-terminal on LHS.
(C) Rules depend on context symbols.
(D) Generates recursively enumerable.

88. If a problem is tractable, it generally means that there exists an algorithm to solve it that runs in:
(A) Polynomial time.
(B) Exponential time.
(C) Logarithmic time.
(D) Constant time.

89. Consider a grammar with the production `expr → expr - expr`. This production rule is:
(A) Right recursive.
(B) Unambiguous.
(C) Normal Form.
(D) Left recursive.

90. An LALR(1) parser is generated by merging states of an LR(1) parser that have the same core set of items. What is a potential consequence of this merging?
(A) It may introduce reduce conflicts.
(B) It increases shift-reduce count.
(C) It parses a larger class range.
(D) It slows down execution time.

91. A compiler designer is creating a system for a new programming language. The phase where the compiler checks if a variable is being used before it has been declared would occur during:
(A) Lexical Analysis
(B) Syntax Analysis
(C) Semantic Analysis
(D) Code Generation

92. Which of the following is a key reason for using an intermediate code representation in a compiler?
(A) Separation of front and back end.
(B) Faster executable code generation.
(C) Better readability for humans.
(D) Removal of the symbol table.

93. A live variable at a certain point in a program is a variable whose value:
(A) Was recently defined.
(B) Stays constant throughout.
(C) May be used before redefine.
(D) Is held in a register.

94. What is the role of the finite automaton in lexical analysis?
(A) Generating parse trees.
(B) Recognizing lexemes.
(C) Checking variable types.
(D) Managing call stacks.

95. In the Chomsky Hierarchy, which type of grammar corresponds to the most powerful class of languages that can be recognized by an algorithm (i.e., Turing-decidable)?
(A) Unrestricted grammars.
(B) Context-sensitive grammars.
(C) Context-free grammars.
(D) Regular grammars.

96. Consider the following statements regarding Pushdown Automata (PDA).
(i) A PDA is essentially a finite automaton augmented with a stack.
(ii) A language is context-free if and only if some PDA recognizes it.
(iii) Deterministic PDAs can recognize all context-free languages.
(iv) A PDA's transition depends on the current state, input symbol, and the symbol on top of the stack.
Which of the above statement(s) is/are correct?
(A) (i), (ii) and (iv) only
(B) (i) and (iii) only
(C) (ii), (iii) and (iv) only
(D) All are correct.

97. The concept of pass-by-value in parameter passing means:
(A) The address is passed.
(B) Function gets a copy.
(C) Changes affect original.
(D) Expression is re-evaluated.

98. Which of the following is an example of a machine-dependent code optimization?
(A) Loop invariant motion.
(B) Constant folding.
(C) Instruction scheduling.
(D) Dead code removal.

Downloaded from www.pdf.tube

Which of the above statement(s) is/are correct?
(A) (i) and (iii) only
(B) (ii) and (iv) only
(C) (ii), (iii), and (iv) only
(D) (i), (ii), and (iv) only

188. A compiler that generates code for a hypothetical machine, which is then simulated, is known as a:
(A) Virtual machine compiler.
(B) Cross compiler system.
(C) Just-in-time compiler.
(D) Native compiler system.

189. A static link or access link in an activation record is necessary for implementing:
(A) Lexical scoping rules.
(B) Function recursion.
(C) Dynamic scoping.
(D) Value passing.

190. What is the primary difference between a syntax tree and a parse tree?
(A) Syntax trees are cleaner.
(B) Syntax is always binary.
(C) Parse tree is for bottom.
(D) No difference exists.

191. Arrange the following computational models in decreasing order of their power (from most powerful to least powerful).
(P) Non-deterministic Finite Automaton (NDFA)
(Q) Deterministic Pushdown Automaton (DPDA)
(R) Turing Machine (TM)
(S) Non-deterministic Pushdown Automaton (NPDA)
Choose the correct sequence.
(A) R, S, Q, P
(B) R, Q, S, P
(C) S, R, Q, P
(D) R, S, P, Q

192. Which of the following tasks is the main purpose of the semantic analyzer?
(A) Lexeme classification.
(B) Parse tree building.
(C) Static rule checks.
(D) Machine code output.

193. The language generated by the grammar S -> 0S1 | ε is:
(A) {0^n 1^n where n ≥ 0}
(B) Equal 0s and 1s set.
(C) All binary strings.
(D) Empty language set.

194. Read the following statements regarding compiler phases.
Statement I : Code optimization is an optional phase in a compiler.
Statement II : All code optimizations reduce the size of the executable file.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

195. Which of the following is the most suitable for managing the symbol table in a compiler for a block-structured language, where variable scopes are nested?
(A) Simple data array.
(B) Hash table stack.
(C) Sequential queue.
(D) Priority queue.

196. An SLR parser is a type of LR parser. What does 'S' in SLR stand for?
(A) Simple.
(B) Standard.
(C) Stack.
(D) State.

NODIA

197. Which of the following is NOT a property under which context-free languages are closed?
(A) Union
(B) Concatenation
(C) Kleene Star
(D) Complementation

198. Match List I with List II regarding complexity classes and language types.
List I: (Concept)
P. Tractable Problem
Q. Intractable Problem
R. Undecidable Problem
S. Regular Language
List II: (Example/Class)
1. Halting Problem
2. Sorting an array

3. The set of strings accepted by a DFA
4. Satisfiability problem (in general)

Choose the correct option from those given below:
(A) P-2, Q-4, R-1, S-3
(B) P-4, Q-2, R-1, S-3
(C) P-2, Q-1, R-4, S-3
(D) P-3, Q-4, R-1, S-2

199. What is a handle in the context of a shift-reduce parser?
(A) Reducible stack portion.
(B) Top stack symbol value.
(C) Lookahead symbol value.
(D) Grammar start symbol.

200. Garbage collection is a mechanism in a run-time system used to:
(A) Reclaim heap memory.
(B) Clear the call stack.
(C) Delete temp files.
(D) Remove dead machine.

201. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The language L = {a^p | p is a prime number} is not regular.
Reason (R) : The Pumping Lemma for regular languages can be used to show that any sufficiently long string in L can be pumped to create a new string with a length that is not a prime number, thus proving L is not regular.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

202. Which of these is a grammar transformation technique used to make a grammar suitable for predictive parsing?
(A) Left factoring.
(B) Strength reduction.
(C) Constant folding.
(D) Dead code removal.

203. The primary job of a linker is to:
(A) Resolve cross-references.
(B) Assemble machine code.
(C) Load file into RAM memory.
(D) Optimize source before compilation.

204. Consider the following statements about Turing Machines.
(i) A Turing machine has an infinite tape.
(ii) A Turing machine's head can move both left and right.
(iii) A Non-deterministic Turing machine is computationally more powerful than a deterministic one.
(iv) The set of languages accepted by Turing machines is the set of recursively enumerable languages.

Which of the above statement(s) is/are correct?
(A) (i), (ii) and (iv) only
(B) (i), (ii) and (iii) only
(C) (ii), (iii) and (iv) only
(D) All are correct.

205. If a grammar is in Chomsky Normal Form (CNF), any derivation of a string of length `n` will require how many steps?
(A) 2n - 1
(B) n
(C) n - 1
(D) n^2

206. Which of the following is an example of a global optimization?
(A) Loop invariant motion.
(B) Local constant fold.
(C) Strength for multiply.
(D) Instruction merging.

207. A language that can be accepted by a Pushdown Automaton is a:
(A) Context-Free Language.
(B) Regular Language.
(C) Context-Sensitive.
(D) Recursively Enumerable.

208. What is the purpose of an attribute grammar?
(A) Semantic associations.
(B) Syntax definitions.
(C) Memory specifications.
(D) Optimization models.

209. In which phase of compilation are errors like missing semicolon or unbalanced braces typically detected?

Click Here to Buy book on Amazon

(A) Syntax Analysis.
(B) Lexical Analysis.
(C) Semantic Analysis.
(D) Code Generation.

210. Which of the following is an incorrect statement about LR parsers?
(A) Bottom-up parser type.
(B) Left-to-right scan.
(C) Handle all CFGs.
(D) Use state stacks.

211. What is the fundamental property of a recursively enumerable language for which a Turing machine exists?
(A) May loop on rejection.
(B) Halts on every input.
(C) Always halts rejects.
(D) Finite tape usage.

212. What problem does left-recursion pose for a top-down parser?
(A) Infinite logic loop.
(B) SR parser conflicts.
(C) Large table size.
(D) Lookahead needs.

213. Which of the following is an incorrect characterization of compiler phases?
(A) Optimization is mandatory.
(B) Lexer creates tokens.
(C) Syntax builds trees.
(D) Semantic checks types.

214. Which sequence of compiler phases is correct?
(A) Lexical -> Syntax -> Semantic
(B) Syntax -> Lexical -> Semantic
(C) Lexical -> Semantic -> Syntax
(D) Semantic -> Syntax -> Lexical

215. Consider the following statements about run-time environments.
(i) Stack allocation is used for data whose size is known at compile time and whose lifetime is nested.
(ii) Heap allocation is used for data that is created and destroyed dynamically at run time.
(iii) The code for procedures is typically stored in the stack.
(iv) Global variables are often allocated in a static data area.
(v) An activation record contains information for a single execution of a procedure.
Which of the above statement(s) is/are correct?
(A) (i), (ii), (iv), and (v) only
(B) (i), (iii), and (v) only
(C) (ii), (iii), and (iv) only
(D) All are correct.

216. A DFA is defined by the tuple {Q, Σ, δ, q0, F}. What does F represent?
(A) Accepting states set.
(B) Transition function.
(C) Set of all states.
(D) Initial start state.

217. The process of evaluating expressions with known constant operands at compile time rather than run time is called:
(A) Constant folding.
(B) Strength reduction.
(C) Subexpression removal.
(D) Dead code removal.

218. Read the following statements regarding the context-free language { a^n b^n c^i | n >= 1, i >= 1 }.
Statement I : The language { a^n b^n c^i | n >= 1, i >= 1 } is context-free.
Statement II : A pushdown automaton can match the count of 'a's with 'b's using its stack, and it does not need to relate them to the count of 'c's.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

219. The 'dynamic link' in an activation record is used to:
(A) Access non-local vars.
(B) Restore caller state.
(C) Point to next code.
(D) Manage heap memory.

220. The primary limitation of a regular grammar is that it cannot:
(A) Represent nested logic.
(B) Use infinite strings.
(C) Generate empty text.
(D) Use multiple symbols.

Downloaded from www.pdf.tube

(D) It implies that all computational problems are solvable.

269. Based on the description, a Universal Turing Machine is conceptually analogous to:
(A) a compiler that translates source code to machine code.
(B) a lexical analyzer that generates a stream of tokens.
(C) a modern stored-program computer that can execute any valid program.
(D) a simple calculator that can only perform one specific function.

270. How does a recursively enumerable language differ from a recursive language?
(A) A recursively enumerable language cannot be recognized by any Turing Machine.
(B) A recursively enumerable language is always context-sensitive, whereas a recursive language is not.
(C) There is no difference; the terms are synonymous in computability theory.
(D) A TM recognizing a recursively enumerable language might not halt on strings outside the language.

Direction : Please read the following source carefully.

The study of computation is not only about what can be computed but also about what cannot. This leads to the domain of unsolvable problems, also known as undecidable problems. A problem is unsolvable if no algorithm, and therefore no Turing Machine, can exist that provides a correct yes/no answer for every possible instance of the problem. A canonical example is the Halting Problem, which asks whether a given Turing Machine will halt on a specific input. Alan Turing proved this problem to be unsolvable. Another famous unsolvable problem is the Post Correspondence Problem (PCP), which involves determining if a sequence of dominoes with strings on both halves can form a matching top and bottom string. The unsolvability of PCP is often used to prove that other problems, particularly problems related to Context-Free Languages, are also unsolvable. The theoretical framework for classifying these languages and their corresponding automata is the Chomsky Hierarchy. This hierarchy organizes languages into four nested levels of increasing complexity: regular, context-free, context-sensitive, and unrestricted (recursively enumerable). Regular languages are recognized by finite automata, context-free by pushdown automata, context-sensitive by linear-bounded automata, and unrestricted grammars generate recursively enumerable languages recognized by Turing Machines. This hierarchy provides a clear map of computational power and expressive capability.

Based on the above source, attempt the following five questions.

271. What is the fundamental characteristic of an unsolvable or undecidable problem?
(A) It can only be solved by a non-deterministic Turing Machine.
(B) It takes an exponential amount of time to find a solution.
(C) No algorithm can exist to solve all instances of the problem correctly.
(D) It is a problem for which only an approximate solution can be found.

272. According to the passage, which language class is associated with Turing Machines in the Chomsky Hierarchy?
(A) Regular languages
(B) Unrestricted languages
(C) Context-Free languages
(D) Context-Sensitive languages

NODIA

273. The Halting Problem, as described in the text, deals with the question of whether a Turing Machine will:
(A) produce the correct output for a given input.
(B) accept a string within a polynomial time frame.
(C) run out of memory on its tape.
(D) halt on a particular input string.

274. How is the Post Correspondence Problem (PCP) often utilized in computability theory?
(A) As a tool to demonstrate the unsolvability of other problems.
(B) As a benchmark for measuring the speed of Turing Machines.
(C) To provide an efficient algorithm for matching strings in context-free grammars.
(D) To define the class of context-sensitive languages.

(Q) LR(0) Parsing
(R) Recursive Descent Parsing
(S) LALR(1) Parsing
Choose the most likely sequence.
(A) R, P, Q, S
(B) Q, P, S, R
(C) R, Q, P, S
(D) P, R, Q, S

233. The primary reason for distinguishing between tractable and intractable problems is to:
(A) Classify usability.
(B) Separate solvability.
(C) Identify FA limits.
(D) Isolate quantum sets.

234. In three-address code, the statement `x = y[i]` would likely be translated into which sequence of operations?
(A) `t1 = i * width; t2 = y + t1; x = *t2`
(B) `x = y + i`
(C) `t1 = y[i]; x = t1`
(D) `t1 = addr(y); t2 = t1 + i; x = *t2`

235. What is the key feature of a context-sensitive language grammar rule `$\alpha \rightarrow \beta$`?
(A) LHS is non-terminal.
(B) Length $|\alpha| \leq |\beta|$.
(C) Length $|\alpha| > |\beta|$.
(D) RHS is terminal.

236. A compiler finds that a variable `v` is assigned a value but this value is never used later in any execution path. Removing the assignment to `v` is an example of:
(A) Dead code removal.
(B) Constant folding.
(C) Strength reduction.
(D) Loop unrolling.

237. Consider the following statements regarding compiler design.
(i) Semantic analysis precedes syntax analysis.
(ii) The symbol table is used to store information about entities like variable names and types.
(iii) L-attributed definitions are a special case of S-attributed definitions.
(iv) A DAG can represent common subexpressions more compactly than a syntax tree.
(v) Lexical analysis is responsible for identifying syntax errors.
Which of the above statement(s) is/are correct?
(A) (i) and (v) only
(B) (ii) and (iv) only
(C) (ii), (iii), and (iv) only
(D) (i), (iv), and (v) only

238. The Myhill-Nerode theorem provides a method to:
(A) Minimize DFA states.
(B) Prove CFL limits.
(C) Convert CFG to CNF.
(D) Test undecidability.

239. What is a dangling else problem in parsing?
(A) Association ambiguity.
(B) Missing 'if' error.
(C) Semantic logic bug.
(D) Else clause removal.

240. Read the following statements regarding recursive and recursively enumerable languages.
Statement I : The class of recursive languages is closed under complementation.
Statement II : The class of recursively enumerable languages is closed under complementation.
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

241. A parser generator tool like YACC (Yet Another Compiler-Compiler) typically produces a(n):
(A) LALR(1) parser.
(B) LL(1) parser system.
(C) Recursive descent.
(D) SLR(1) parser system.

242. What is the fundamental mechanism used by a pushdown automaton that allows it to recognize context-free languages but not regular languages?
(A) Stack memory usage.
(B) Non-det transitions.
(C) Multiple final state.
(D) Right-to-left scan.

243. The problem of determining if an arbitrary Turing machine T will ever write a specific symbol 's' on its tape is:

Click Here to Buy book on Amazon

(A) Undecidable.
(B) Decidable.
(C) Tractable.
(D) NP-class.

244. A syntax-directed definition is L-attributed if each attribute of a production's right-hand side symbol Xj is dependent only on:
(A) Parent/Left siblings.
(B) Right sibling data.
(C) Synthesized values.
(D) Parent node values.

245. The intermediate representation that is structured like a graph and is particularly good for identifying common subexpressions is a:
(A) Directed Acyclic Graph (DAG).
(B) 3-address code list.
(C) Postfix string text.
(D) Stack machine code.

246. In the theory of computation, what does Russell's Paradox demonstrate?
(A) Set theory flaw.
(B) Automata limits.
(C) Halting problem.
(D) Tractability gap.

247. Match List I with List II regarding compiler tasks and associated concepts.
List I: (Compiler Task)
P. Parsing
Q. Semantic Analysis
R. Runtime Memory Management
S. Lexical Analysis
List II: (Associated Data Structure/Concept)
1. Stack and Heap
2. Regular Expressions
3. Context-Free Grammar
4. Attribute Grammar
Choose the correct option from those given below:
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-2, S-1
(C) P-3, Q-2, R-1, S-4
(D) P-1, Q-4, R-3, S-2

248. The control flow of a program refers to:
(A) Execution order.
(B) Data flow paths.
(C) Memory deallocation.
(D) Type verification.

249. A context-free grammar G is ambiguous if:
(A) Multiple derivations.
(B) Infinite language.
(C) Useless productions.
(D) Non-regular language.

250. Which of the following is considered the most powerful class of grammar in the Chomsky hierarchy?
(A) Type-0 Grammar
(B) Type-3 Grammar
(C) Type-2 Grammar
(D) Type-1 Grammar

NODIA

Direction : Please read the following source carefully.

The foundational class of languages in formal language theory is the regular languages, which are characterized by a variety of equivalent models of computation and representation. The most fundamental computational model is the Deterministic Finite Automaton (DFA), a machine that reads an input string one symbol at a time and exists in one of a finite number of states, transitioning between them based on the current state and the input symbol. A variation is the Non-Deterministic Finite Automaton (NDFA), which can have multiple possible next states for a given state-symbol pair and can include transitions on an empty string. Despite its non-determinism, a key theorem in automata theory establishes that for every NDFA, there is an equivalent DFA that recognizes the same language. This equivalence is pivotal, as it allows for the more flexible design of NDFAs while guaranteeing that the underlying language is still regular and can be implemented deterministically. Regular languages can also be described declaratively using regular grammars, which have restricted production rules, or through regular expressions, a concise algebraic notation for specifying string patterns. The equivalence among DFAs, NDFAs, regular grammars, and regular expressions forms the cornerstone of the theory of regular languages, providing a robust framework for their study and application.

Based on the above source, attempt the following five questions.

251. According to the passage, what is the primary significance of the equivalence between DFAs and NDFAs?
(A) It proves that NDFAs are computationally more powerful than DFAs.
(B) It allows for flexible automaton design with the assurance of deterministic implementability.
(C) It demonstrates that non-determinism adds a new class of languages to formal theory.
(D) It establishes that regular grammars are superior to regular expressions for language definition.

252. Which of the following is NOT mentioned in the passage as a model for representing regular languages?
(A) Pushdown Automaton (PDA)
(B) Deterministic Finite Automaton (DFA)
(C) Regular Expressions
(D) Non-Deterministic Finite Automaton (NDFA)

253. Based on the description in the passage, a key difference between a DFA and an NDFA is that an NDFA can:
(A) have an infinite number of states.
(B) process input strings from right to left.
(C) have multiple transition possibilities for a single input symbol.
(D) recognize languages that a DFA cannot.

254. The passage suggests that the various models for regular languages are important because they:
(A) create a hierarchy of computational power.
(B) are primarily used for generating non-computational problems.
(C) each define a completely different and separate set of languages.
(D) provide a robust and interconnected theoretical framework.

255. What can be inferred about the relationship between regular grammars and regular expressions from the text?
(A) Regular grammars are a procedural model, while regular expressions are a declarative notation.
(B) Regular expressions are used to build DFAs, whereas regular grammars are used to build NDFAs.
(C) They are two distinct and equivalent methods for describing the same class of languages.
(D) Regular grammars have more restrictive production rules than regular expressions have operators.

Direction : Please read the following source carefully.

The class of regular languages exhibits several important properties and has a fundamental limitation that defines its boundary. These languages are closed under operations such as union, intersection, concatenation, and complementation, meaning that if one or more languages are regular, the result of applying these operations will also be a regular language. This closure is critical for building complex regular language recognizers from simpler ones. However, not all languages are regular. The primary tool for proving that a language is non-regular is the Pumping Lemma for regular languages. This lemma states that for any regular language, any sufficiently long string in the language can be pumped – that is, a middle section of the string can be repeated any number of times (including zero) to produce a new string that is also in the language. To prove a language is non-regular, one can use a proof by contradiction: assume the language is regular, apply the Pumping Lemma to a chosen string, and show that pumping it leads to a string that is not in the language, thus violating the lemma. This theoretical framework has a direct practical application in the first phase of a compiler, known as lexical analysis. A lexical analyzer, or scanner, uses the principles of regular languages, typically specified by regular expressions, to partition the source code into a stream of tokens.

Based on the above source, attempt the following five questions.

256. Which concept is described in the passage as the primary method for demonstrating that a language does not belong to the regular class?
(A) The closure properties of regular languages.
(B) The application of the Pumping Lemma.
(C) The process of lexical analysis.
(D) The equivalence of DFAs and NDFAs.

Downloaded from www.pdf.tube

221. Match List I with List II regarding complexity classes and their characteristics.
List I: (Problem Class)
P. Class P
Q. Class NP
R. Undecidable
S. NP-Hard
List II: (Characteristic)
1. No algorithm exists to solve it.
2. Solvable in deterministic polynomial time.
3. At least as hard as any problem in NP.
4. Solution can be verified in polynomial time.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-4, R-1, S-3
(C) P-2, Q-3, R-1, S-4
(D) P-4, Q-2, R-3, S-1

222. What is the main characteristic of an S-attributed syntax-directed definition?
(A) Synthesized attributes.
(B) Only inherited attributes.
(C) Both attribute types.
(D) Top-down evaluation.

223. What is the output of the lexical analysis phase of a compiler?
(A) A stream of tokens.
(B) A parse tree model.
(C) Abstract syntax tree.
(D) Machine code output.

224. The equivalence of NFAs and DFAs implies that for any NFA:
(A) Equivalent DFA exists.
(B) Larger DFA exists.
(C) Same-state DFA exists.
(D) No DFA is possible.

225. A thunk is a concept most closely associated with which parameter-passing mechanism?
(A) Pass-by-name.
(B) Pass-by-value.
(C) Pass-by-reference.
(D) Pass-by-result.

226. What does a context-free grammar consist of?
(A) Terminals and rules.
(B) States and transitions.
(C) Tape and machine head.
(D) Stack and control unit.

227. Given below are two statements, one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A Universal Turing Machine can compute any function that is computable.
Reason (R) : It operates by taking a description of another Turing Machine and its input, and then simulating the behavior of that machine on that input.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

228. The action of a parser to use a grammar rule to replace a sequence of symbols on the stack with a single non-terminal is called:
(A) Reduce.
(B) Shift.
(C) Accept.
(D) Error.

229. Which of the following is NOT a typical component of an activation record?
(A) Source code text.
(B) Local data values.
(C) Return address.
(D) Function parameters.

230. A loop optimization technique that moves a computation from inside a loop to outside the loop is called:
(A) Loop-invariant motion.
(B) Loop body unrolling.
(C) Nested loop fusion.
(D) Nested loop fission.

231. The language L = {a^i b^j c^k | i = j or j = k} is:
(A) Context-free.
(B) Regular type.
(C) Not context-free.
(D) Both reg and CF.

NODIA

232. Arrange the following parsing algorithms in chronological order of their conceptual development.
(P) LL(1) Parsing

263. The passage states that Non-Deterministic Pushdown Automata (NPDAs) are equivalent in power to which of the following?
(A) Deterministic Finite Automata (DFAs)
(B) Turing Machines (TMs)
(C) Context-Free Grammars (CFGs)
(D) Regular Grammars

264. What is the stated purpose of converting a Context-Free Grammar into Chomsky Normal Form or Greibach Normal Form?
(A) To eliminate ambiguity from the grammar.
(B) To increase the expressive power of the grammar.
(C) To prove that the language generated is non-regular.
(D) For standardization and to improve parsing efficiency.

265. What can be inferred about the relationship between deterministic and non-deterministic pushdown automata from the passage?
(A) Non-deterministic PDAs can recognize a broader class of languages than deterministic ones.
(B) Deterministic PDAs are a theoretical concept with no practical application.
(C) They are equivalent in their language recognition power.
(D) Every NPDA has an equivalent DPDA, similar to the NFA-DFA relationship.

Direction : Please read the following source carefully.

At the apex of automata theory sits the Turing Machine (TM), a powerful model of computation that serves as the theoretical basis for modern computing. A standard Turing Machine consists of a finite control unit, a tape that is infinite in one or both directions, and a read/write head that can move along the tape. This model, despite its simplicity, can simulate any computer algorithm. Variations of the standard TM, such as multi-tape or non-deterministic TMs, have been proposed, but it has been proven that they are all equivalent in computational power to the standard TM, as they can be simulated by it. This robustness leads to the Church-Turing Thesis, a fundamental hypothesis which posits that any function that can be computed by an algorithm can be computed by a Turing Machine. TMs define the class of Recursive and Recursively Enumerable Languages. A language is recursively enumerable if a TM can halt and accept any string in the language, but may loop forever on strings not in the language. A language is recursive if a TM can halt on all inputs, definitively accepting or rejecting each string. The ability to construct a TM for a problem, even a simple one, is a demonstration of its computability. A Universal Turing Machine is a specific TM designed to simulate any other Turing Machine, embodying the principle of the stored-program computer.

Based on the above source, attempt the following five questions.

266. What does the Church-Turing Thesis propose?
(A) That all computational problems can be solved in a finite amount of time.
(B) That any function computable by an algorithmic process can be computed by a Turing Machine.
(C) That Turing Machines are computationally equivalent to Pushdown Automata.
(D) That non-deterministic Turing Machines are more powerful than standard Turing Machines.

267. According to the passage, what is the defining characteristic of a recursive language?
(A) A Turing Machine will halt and accept strings in the language but may loop on others.
(B) A Turing Machine exists that halts on all possible input strings, either accepting or rejecting them.
(C) The language can be generated by a context-free grammar.
(D) It can be recognized by a Universal Turing Machine exclusively.

268. What is the significance of the fact that variations of the Turing Machine are equivalent in power to the standard model?
(A) It highlights the robustness of the Turing Machine as a universal model of computation.
(B) It demonstrates the inherent limitations of computational models.
(C) It proves that non-determinism does not add computational power at this level.

Click Here to Buy book on Amazon

257. According to the passage, if L1 and L2 are regular languages, which of the following operations is guaranteed to produce another regular language?
(A) The subdivision of L1 into smaller components.
(B) The mapping of L1 to a context-free language.
(C) The pumping of a string that belongs to L1 but not L2.
(D) The intersection of L1 and L2.

258. What is the role of lexical analysis as described in the text?
(A) To use regular language principles to identify tokens in source code.
(B) To verify the grammatical structure of the entire program.
(C) To prove that the source code's language is non-regular.
(D) To apply the Pumping Lemma to optimize the program's execution.

259. How is the Pumping Lemma for regular languages typically used?
(A) To construct a DFA for a given regular language.
(B) To show that a language is regular by finding a string that can be pumped.
(C) To prove a language is non-regular by showing a contradiction.
(D) To convert a regular expression into a regular grammar.

260. What can be inferred about the strings in a regular language from the description of the Pumping Lemma?
(A) Sufficiently long strings must have a repetitive structure that allows for pumping.
(B) Only short strings can belong to a regular language.
(C) All strings in a regular language must contain a repeated section.
(D) Pumping a string always makes it longer.

Direction : Please read the following source carefully.

Moving up the Chomsky hierarchy from regular languages, we encounter the Context-Free Languages (CFLs), which can model more complex structures like nested expressions common in programming languages. The primary generative mechanism for these languages is the Context-Free Grammar (CFG), a set of production rules where the left-hand side is a single non-terminal symbol. The computational model that recognizes the class of CFLs is the Pushdown Automaton (PDA). A PDA is essentially a Non-Deterministic Finite Automaton (NDFA) augmented with a stack, which provides it with unbounded memory. This stack allows the PDA to recognize languages that require matching symbols, such as balanced parentheses, which finite automata cannot. While PDAs can be deterministic, the Non-Deterministic Pushdown Automaton (NPDA) is more powerful and is equivalent in recognition power to CFGs. For standardization and parsing efficiency, CFGs can be converted into normal forms, such as Chomsky Normal Form (CNF) and Greibach Normal Form (GNF). A single derivation in a CFG can be visually represented by a parse tree. However, if a grammar can produce more than one parse tree for the same string, the grammar is said to possess ambiguity, a generally undesirable property in programming language design. The equivalence of NPDAs and CFGs is a cornerstone of this language class, similar to the DFA/NDFA equivalence for regular languages.

Based on the above source, attempt the following five questions.

261. What key component distinguishes a Pushdown Automaton from a Finite Automaton, enabling it to recognize Context-Free Languages?
(A) A read-write tape
(B) A stack
(C) Multiple start states
(D) A stack

262. According to the passage, what does it mean for a Context-Free Grammar to be ambiguous?
(A) The grammar generates a non-regular language.
(B) A string in the language can have more than one corresponding parse tree.
(C) The grammar cannot be converted into Chomsky Normal Form.
(D) The equivalent Pushdown Automaton for the grammar must be deterministic.

275. Based on the Chomsky Hierarchy described in the passage, which statement is true?
(A) All context-free languages are also regular languages.
(B) Pushdown automata are more powerful than Turing Machines.
(C) The class of regular languages is a subset of the class of context-free languages.
(D) Context-sensitive languages are generated by unrestricted grammars.

Direction : Please read the following source carefully.

Syntax analysis, or parsing, is the second phase of compilation, where the linear sequence of tokens from the lexical analyzer is structured into a hierarchical representation like a parse tree. This process verifies if the token stream conforms to the syntactic rules defined by a grammar. Parsers are broadly categorized into two types: top-down and bottom-up. Top-down parsing, such as Recursive Descent or LL(1) parsing, attempts to construct the parse tree from the root downwards. It works by starting with the grammar's start symbol and applying productions to derive the source string. In contrast, bottom-up parsing, exemplified by LR parsers (and their variant LALR(1)), works by reducing the input string back to the start symbol, effectively building the parse tree from the leaves upwards. A significant challenge in parser construction is handling grammar ambiguities and conflicts. To resolve these, language specifications often rely on rules of associativity (e.g., left-associativity for subtraction) and precedence (e.g., multiplication before addition). When these rules are insufficient, grammar transformations may be required to rewrite the grammar into an equivalent form that is suitable for a specific parsing method, such as eliminating left recursion for top-down parsers. These techniques are crucial for building efficient and deterministic parsers for programming languages.

Based on the above source, attempt the following five questions.

276. What is the fundamental difference between top-down and bottom-up parsing?
(A) Top-down parsers build the parse tree from the root, while bottom-up parsers build it from the leaves.
(B) Top-down parsers use a stack, whereas bottom-up parsers do not.
(C) Top-down parsers are used for regular languages, while bottom-up parsers are for context-free languages.
(D) Top-down parsers handle precedence and associativity, while bottom-up parsers require grammar transformations.

277. According to the passage, what is the role of precedence and associativity rules in syntax analysis?
(A) To transform the grammar into an LL(1) format.
(B) To convert the token stream into an abstract syntax tree.
(C) To define the data types of variables and expressions.
(D) To resolve ambiguities in how expressions should be structured.

NODIA

278. Which of the following is mentioned as an example of a bottom-up parsing method?
(A) Recursive Descent Parsing
(B) LR Parser
(C) LL(1) Parsing
(D) Predictive Parsing

279. The passage suggests that grammar transformations, such as eliminating left recursion, are often necessary for what purpose?
(A) To increase the number of tokens generated by the lexical analyzer.
(B) To prove that the language is context-free.
(C) To make a grammar suitable for a top-down parsing algorithm.
(D) To convert a bottom-up parser into a top-down parser.

280. What can be inferred about the input and output of the syntax analysis phase?
(A) It takes source code as input and produces machine code as output.
(B) It takes a stream of characters as input and produces a stream of tokens as output.
(C) It takes a parse tree as input and produces an optimized parse tree as output.
(D) It takes a stream of tokens as input and produces a hierarchical syntactic structure as output.

Direction : Please read the following source carefully.

Following successful syntax analysis, a compiler proceeds to semantic analysis, a phase dedicated to checking the source program for semantic errors and gathering type information for the subsequent code generation phase. This phase operates on the hierarchical structure, typically a parse tree or an abstract syntax tree, produced by the parser. A formal framework for specifying semantic rules is the attribute grammar, which associates semantic attributes with grammar symbols. These attributes, which can represent information like data types, memory locations, or values, are computed by semantic rules. Attributes are categorized as either synthesized or inherited. Synthesized attributes are computed from the values of attributes at the children nodes in a parse tree, effectively passing information up the tree. In contrast, inherited attributes are computed from the attributes at the parent and/or sibling nodes, passing information down or across the tree. A Syntax Directed Definition (SDD) is a context-free grammar augmented with semantic rules. The dependencies between attributes can be visualized using a dependency graph, which helps determine the evaluation order for the attributes. SDDs are classified as S-attributed if they only use synthesized attributes, and L-attributed if attributes can be inherited, but only from left siblings. A primary function of this phase is type-checking, which ensures that operators are applied to compatible operands.

Based on the above source, attempt the following five questions.

281. What is the primary function of semantic analysis in a compiler?
(A) To check the program for semantic correctness and perform type-checking.
(B) To parse the source code into a hierarchical structure.
(C) To generate optimized, low-level machine code.
(D) To break the source program into a stream of tokens.

282. How is a synthesized attribute's value determined?
(A) It is passed down from the parent node in the parse tree.
(B) It is computed from the attributes of its children nodes.
(C) It is determined by attributes of its left siblings only.
(D) It is defined globally and is constant throughout the tree.

283. According to the passage, what does a dependency graph represent?
(A) The flow of control within the source program.
(B) The relationships and evaluation order among attributes.
(C) The hierarchy of function calls and procedure activations.
(D) The syntactic structure of the program according to the grammar.

284. An S-attributed definition is a specific type of Syntax Directed Definition that:
(A) uses only synthesized attributes.
(B) allows both synthesized and inherited attributes without restriction.
(C) uses only inherited attributes.
(D) is used exclusively for type-checking.

285. What is the relationship between an attribute grammar and a Syntax Directed Definition (SDD)?
(A) An attribute grammar is a visual representation of an SDD.
(B) An attribute grammar is used for syntax analysis, while an SDD is used for semantic analysis.
(C) An SDD is a grammar augmented with semantic rules, which is a formal way to express an attribute grammar.
(D) An SDD is a less formal specification than an attribute grammar.

Direction : Please read the following source carefully.

The execution of a program requires a run-time system to manage memory and program execution flow. A crucial aspect of this system is the storage organization, which dictates how memory is allocated for program data and control information. For procedural languages, the execution flow is often visualized as an activation tree, where each node represents an activation of a procedure. The information needed for a single

Downloaded from www.pdf.tube

procedure activation is stored in an activation record (or stack frame). A typical activation record contains space for local variables, temporary values, saved machine status (like the program counter), an access link to non-local data, and the actual parameters passed to the procedure. The most common strategy for managing these records is stack allocation, where activation records are pushed onto a run-time stack when a procedure is called and popped off when it returns. This LIFO (Last-In, First-Out) behavior perfectly mirrors the nested, sequential nature of procedure calls and returns. The way arguments are provided to a procedure is governed by parameter passing mechanisms, such as pass-by-value or pass-by-reference. Throughout this process, the symbol table, a data structure maintained by the compiler, is used to store information about identifiers like variables and functions, which is then used by the run-time system to manage storage.

Based on the above source, attempt the following five questions.

286. In the context of a run-time system, what does an activation record typically contain?
(A) The entire source code of the called procedure.
(B) Only the return address and the actual parameters.
(C) A pointer to the symbol table and the lexical analyzer's output.
(D) Information like local variables, parameters, and saved machine status.

287. What does the stack allocation strategy for activation records directly correspond to?
(A) The order of variable declarations in the source code.
(B) The physical layout of the computer's hard disk.
(C) The nested, LIFO nature of procedure calls and returns.
(D) The complexity of the algorithm implemented by the procedure.

288. According to the passage, the activation tree is a visualization of what aspect of a program's execution?
(A) The flow of data between different modules.
(B) The sequence of procedure activations.
(C) The memory addresses allocated to variables.
(D) The dependencies between semantic attributes.

289. The passage mentions that the symbol table is used by the run-time system for what purpose?
(A) To perform type-checking on procedure arguments.
(B) To generate intermediate code from the abstract syntax tree.
(C) To optimize the program by reordering instructions.
(D) To manage storage based on information about identifiers.

290. Which of the following is described as a component managed by the run-time system?
(A) The specific mechanisms used for passing parameters.
(B) The parsing of the source program into a syntax tree.
(C) The conversion of a grammar to Chomsky Normal Form.
(D) The lexical analysis of the source code into tokens.

Direction : Please read the following source carefully.

After the source program has been validated by the syntax and semantic analysis phases, many compilers translate it into an explicit low-level or machine-like intermediate representation (IR). This intermediate code generation phase acts as a bridge between the front end (analysis) and the back end (synthesis) of the compiler. The IR is designed to be machine-independent, making it easier to retarget the compiler to different machine architectures, and it is also in a form that is amenable to optimization. There are various forms of intermediate representations, such as three-address code, syntax trees, and stack-based code. The translation process involves systematically converting high-level programming constructs into this IR. For instance, declarations are translated into symbol table entries with associated type and storage information. Assignment statements are typically broken down into a sequence of simpler operations. Control flow constructs like if-then-else statements and while loops are translated using labels and conditional or unconditional jumps. Boolean expressions are

NODIA

ANSWERS

1. (B) A subset of strings of symbols derived from a specific alphabet.
In computation theory, a formal language is a set of strings composed of symbols taken from a fixed alphabet. These strings are formed according to specific rules, often defined by grammars. This concept allows for the mathematical study of syntax and structures within both natural and artificial systems, serving as the basis for automata theory.

2. (A) That some problems are computationally unsolvable.
Georg Cantor's diagonal argument is a pivotal proof technique used to demonstrate that certain infinite sets, such as real numbers, are larger than others, like natural numbers. In computation, it proves the existence of undecidable problems by showing that no algorithm can map all possible inputs to correct outputs, thus establishing fundamental limits.

3. (C) If no algorithm can solve the problem for all possible inputs.
Non-computational problems are those for which no systematic procedure or algorithm can be devised to provide a correct solution for every instance in a finite amount of time. Unlike solvable problems that might simply require more resources, these are theoretically impossible to resolve computationally, regardless of the hardware power or processing time available.

4. (A) Both Statement I and Statement II are true.
Russell's Paradox identifies a logical contradiction in naive set theory by considering a set containing all sets that do not contain themselves. If such a set exists, it must both contain and not contain itself. This discovery necessitated the development of axiomatic set theory, such as Zermelo-Fraenkel, to provide a consistent logical foundation.

5. (C) A supplementary stack memory unit.
A deterministic finite automaton is formally defined by five components: a finite set of states, an input alphabet, a transition function, a start state, and a set of final states. While it processes input strings through these transitions, it lacks any auxiliary memory structures like a stack or tape, which are characteristics of more powerful machines.

6. (A) P-4, Q-3, R-1, S-2
The theory of computation links specific proof techniques and paradoxes to language definitions and problem classes. The diagonal argument proves differing cardinalities, while Russell's Paradox explores set membership contradictions. Formal languages are defined as string subsets over alphabets, and non-computational problems are categorized by the total absence of a resolving algorithm for all possible inputs.

7. (D) A set of states.
In a non-deterministic finite automaton, the transition function maps a state and an input symbol to a subset of the state set. This allows the machine to be in multiple states simultaneously or choose between different paths for the same input. Unlike deterministic models, this flexibility includes transitions without consuming input, effectively mapping to a power set.

8. (A) Every language accepted by an NDFA is also accepted by a DFA.
Although non-deterministic machines appear more flexible, they are equivalent in expressive power to deterministic finite automata. Every language that can be recognized by an NDFA can also be recognized by some DFA. This equivalence is proven through the subset construction algorithm, which translates non-deterministic transitions into a single deterministic path across a larger set of states.

9. (A) Both A and R are true and R is the correct explanation of A.
The class of regular languages is precisely defined by what both deterministic and non-deterministic finite automata can accept. Since a deterministic machine is technically a restricted version of a non-deterministic one, and any non-deterministic machine can be transformed into a deterministic equivalent via the subset

Click Here to Buy book on Amazon

www.nodia.in

intermediate code.
(B) To convert the intermediate code into a human-readable format.
(C) To translate the source code into an intermediate representation.
(D) To improve the run-time performance or memory usage of the program.

297. How does peep-hole optimization work?
(A) By examining a small, moving sequence of instructions and making local improvements.
(B) By analyzing the entire control-flow graph of the program.
(C) By optimizing only the code within loops.
(D) By re-allocating all registers for an entire function.

298. According to the passage, data-flow analysis is essential for what purpose?
(A) Building the initial parse tree from the token stream.
(B) Understanding how data values are propagated through the program for optimization.
(C) Allocating memory on the run-time stack for activation records.
(D) Checking for type mismatches between operators and operands.

299. What distinguishes global optimization from local optimization?
(A) Global optimization focuses on loops, while local optimization focuses on function calls.
(B) Global optimization operates on larger program units like functions, while local optimization works on basic blocks.
(C) Global optimization is performed on the source code, while local is on the target code.
(D) Global optimization is machine-dependent, while local optimization is machine-independent.

NODIA

300. The final step of instruction scheduling, as mentioned in the text, is part of which compiler phase?
(A) Code Generation
(B) Intermediate Code Generation
(C) Semantic Analysis
(D) Syntax Analysis

symbol with a marker, move across the tape to find and mark a corresponding instance of the second symbol, and repeat this cycle until all symbols are correctly matched.

37. (C) The machine M eventually halts its computation.
The Halting Problem is a fundamental question in computation that asks whether an arbitrary program, modeled as a Turing Machine, will eventually stop running or continue forever when given a specific input. Alan Turing proved that this problem is undecidable, meaning it is impossible to create a general algorithm that can determine the halting status for all possible machine-input pairs.

38. (A) Finding string sequences with matching concatenations.
The Post Correspondence Problem is a classic undecidable problem in formal language theory. It involves two sets of strings and asks whether there is a sequence of indices such that concatenating the strings from the first set in that order results in the same string as concatenating the corresponding strings from the second set. Its undecidability demonstrates string manipulation limits.

39. (C) Checking if a given CFG is ambiguous.
While several properties of context-free languages are decidable, determining whether a given context-free grammar is ambiguous is famously undecidable. There is no general algorithm that can examine any context-free grammar and correctly decide if it can produce more than one parse tree for some string. This limitation poses a significant challenge for developers creating tools that generate parsers.

40. (A) Both Statement I and Statement II are true.
Tractability is a measure of how efficiently a problem can be solved by a computer. Problems with algorithms that run in polynomial time are considered tractable because they are generally solvable for large inputs. In contrast, problems requiring super-polynomial time, such as exponential time, are considered intractable as they quickly become impossible to solve practically as the input size increases.

41. (B) Evaluation order.
During the parsing phase of compilation, associativity and precedence rules are essential for resolving ambiguities in expressions. These rules define the specific order in which operators are evaluated when multiple operators appear in the same expression. Precedence determines which operations happen first, while associativity handles the evaluation direction for operators of the same level, ensuring consistent and correct program logic.

42. (A) $E \rightarrow E + T \mid T, T \rightarrow T * F \mid F, F \rightarrow id$
To resolve ambiguity and enforce standard mathematical order in a grammar, it must be restructured to reflect different levels of precedence. This transformation typically involves introducing new non-terminals for each level, such as terms and factors. By nesting these rules, the grammar ensures that operations like multiplication are performed before addition, creating a unique and correct parse tree.

43. (A) (i), (ii) and (iv) only
Top-down parsing is a strategy that constructs a parse tree starting from the root symbol and moving toward the leaves. It generates a leftmost derivation of the input string as it processes tokens. A significant challenge for this method is left-recursion in the grammar, which can lead to infinite loops. Unlike bottom-up approaches, it typically scans input from left to right.

44. (B) Uses lookahead to choose rules.
A predictive parser, a type of top-down parser, avoids the need for backtracking by using a lookahead mechanism. By examining the next one or more symbols in the input stream, the parser can uniquely determine which grammar production rule to apply. This efficiency allows the parser to operate in linear time, provided the grammar is carefully designed to avoid conflicts.

45. (A) Mutually recursive procedures for non-terminals.
Recursive descent parsing is a top-down technique where each non-terminal in the grammar is represented by a specific procedure or function. These procedures call each other, mirroring the structure of the grammar rules. This approach relies on recursion to handle nested components, allowing the parser to construct a parse tree by

Downloaded from www.pdf.tube

construction method, both models recognize exactly the same set of languages.

10. (A) a(a+b)*b
A regular expression for strings starting with one character and ending with another must explicitly define those boundaries. For strings over an alphabet that begin with a specific symbol and terminate with another, any combination of symbols from the alphabet can appear in between. This structure ensures the prefix and suffix requirements are met for every string.

11. (D) A → aB or A → a
Regular grammars are restricted to specific production formats to generate regular languages. In a right-linear grammar, rules must either transition a non-terminal to a terminal followed by another non-terminal or simply to a terminal symbol. This specific structure ensures that the generated strings can be recognized by a finite automaton without needing any auxiliary stack memory.

12. (D) (i), (ii), (iii), and (iv)
Regular languages exhibit strong closure properties under various operations. They remain regular when subjected to union, intersection, complementation, and concatenation. This means that performing these mathematical operations on regular sets always results in a set that can still be recognized by a finite automaton. Such robustness allows for complex language constructions while maintaining the simplicity of the underlying model.

13. (C) Prove that a given language is not regular.
The Pumping Lemma serves as a necessary condition for a language to be classified as regular. It states that all sufficiently long strings in a regular language can be divided and pumped by repeating a middle section. By demonstrating that a language lacks this property, researchers can rigorously prove that the language cannot be recognized by any finite automaton.

14. (D) It requires counting symbols beyond finite memory.
Finite automata possess limited memory, represented by a fixed number of states. A language requiring the comparison of counts for two different symbols, such as an equal number of two distinct characters, requires potentially infinite memory to track the count of the first symbol. Since a finite machine cannot perform this unbounded counting, the language is considered non-regular.

15. (A) To group characters into meaningful lexemes and produce tokens.
Lexical analysis is the initial phase of a compiler where the source code characters are scanned and grouped into meaningful sequences called lexemes. These lexemes are then categorized into tokens, which represent the basic building blocks of the language syntax. This process simplifies the subsequent parsing stage by filtering out unnecessary elements like comments and excessive whitespace.

NODIA

16. (B) A stack
A pushdown automaton extends the capabilities of a finite automaton by incorporating an external stack. This last-in, first-out data structure allows the machine to store and retrieve symbols, enabling it to recognize context-free languages. The stack provides the necessary memory to handle nested structures and balanced symbols, which are common requirements in programming language syntax and logic.

17. (A) R, P, S, Q
The Chomsky Hierarchy categorizes grammars and languages into four levels based on their expressive power. Regular languages are the least expressive, followed by context-free languages, which introduce nested structures. Context-sensitive languages add further constraints based on surrounding symbols, while recursively enumerable languages are the most powerful, representing any language that can be recognized by a standard Turing machine.

18. (A) NPDAs recognize a larger class of languages than DPDAs can.
While deterministic and non-deterministic finite automata are equivalent in power, the same does not hold for pushdown automata. Non-deterministic pushdown automata can recognize a broader class of context-free languages than their deterministic counterparts. This disparity arises because certain context-free structures require non-deterministic choices to resolve

often handled by creating a sequence of jumps or by using boolean values directly. Similarly, procedure calls are translated into a sequence of actions that include setting up the activation record, passing parameters, and handling the return value. This methodical translation ensures that the program's logic is preserved in a form suitable for subsequent processing.

Based on the above source, attempt the following five questions.

291. What is a primary advantage of using a machine-independent intermediate representation (IR)?
(A) It executes faster than the original high-level source code.
(B) It eliminates the need for semantic analysis.
(C) It allows the compiler to be more easily adapted to different target machines.
(D) It is directly understandable by human programmers.

292. How are high-level control flow statements like 'if-then-else' typically translated into intermediate code?
(A) By using a combination of labels and conditional jump instructions.
(B) By converting them into a single, complex instruction.
(C) By creating a separate function for each branch of the statement.
(D) By ignoring them, as control flow is handled by the target machine.

293. According to the passage, the intermediate code generation phase connects which two major parts of a compiler?
(A) The lexical analyzer and the syntax analyzer.
(B) The symbol table manager and the error handler.
(C) The code optimizer and the target code generator.
(D) The front-end analysis and the back-end synthesis.

294. Which of the following is NOT mentioned in the text as a form of intermediate representation?
(A) Three-address code
(B) Source code
(C) Syntax trees
(D) Stack-based code

295. The translation of a procedure call into intermediate code involves which of the following actions?
(A) Re-parsing the entire program to find the procedure definition.
(B) Translating the procedure's body into machine code immediately.
(C) Generating code to set up an activation record and pass parameters.
(D) Optimizing the procedure's logic using peep-hole optimization.

Direction : Please read the following source carefully.

The final phases of a compiler are code generation and code optimization, which are responsible for producing efficient target code from the intermediate representation. Code optimization aims to improve the intermediate code to produce a program that runs faster or uses less memory. This is not a single step but a collection of techniques. A prerequisite for many advanced optimizations is the analysis of the program's structure, specifically its control-flow and data-flow properties. Control-flow analysis builds a graph representing the possible execution paths, while data-flow analysis gathers information about how data values are transmitted through the program. Optimization techniques can be classified by their scope. Local optimization works on basic blocks (straight-line code sequences), while global optimization operates across entire functions or procedures. Loop optimization is particularly important as programs spend much of their time in loops; techniques include code motion and induction variable elimination. A simpler but effective technique is peep-hole optimization, which examines a small, sliding window of target code instructions and replaces them with a shorter or faster sequence. Finally, the code generation phase translates the optimized IR into the target machine's instruction set, which involves instruction selection, register allocation, and instruction scheduling to reorder instructions for optimal performance on modern processors.

Based on the above source, attempt the following five questions.

296. What is the primary purpose of code optimization?
(A) To verify the syntactic correctness of the

Click Here to Buy book on Amazon

can be generated by a context-free grammar. This duality is fundamental to language study.

28. (B) Every computable function is computable by a TM.
The Church-Turing Thesis is a fundamental principle in computer science which asserts that any function that can be calculated by an algorithm can also be computed by a Turing machine. It effectively defines the limits of what is computable, suggesting that the Turing machine is a universal model encompassing all possible effective methods for solving mathematical and logical problems.

29. (A) Both A and R are true and R is the correct explanation of A.
A Universal Turing Machine is a versatile model capable of simulating any other Turing machine. It achieves this by taking the encoded description of a specific machine and an input string as its own input. The UTM then follows the encoded instructions to replicate the behavior of the described machine, demonstrating that a single general-purpose machine can perform any computation.

30. (D) A Turing Machine halts on all inputs to accept/reject.
A language is classified as recursive if there exists a Turing machine that decides it, meaning the machine halts and provides a definite accept or reject answer for every possible input string. Unlike recursively enumerable languages where a machine might loop forever on rejected inputs, recursive languages ensure that the computation always terminates, making them fully decidable by algorithms.

31. (C) (i), (ii), (iii), and (iv)
The Chomsky Hierarchy establishes a nested relationship between different language classes. Regular languages are a subset of context-free languages, which are themselves contained within context-sensitive languages. Furthermore, context-sensitive languages are a subset of recursive languages, and the set of recursive languages is a subset of recursively enumerable languages. This hierarchical structure defines increasing expressive power.

32. (B) R is a proper subset of RE.
In the theory of computation, recursive languages represent a stricter category than recursively enumerable languages. While every recursive language is also recursively enumerable, there are recursively enumerable languages that are not recursive. This means that the set of recursive languages is a proper subset of the set of recursively enumerable languages, highlighting the distinction between semi-decidability and full decidability.

33. (B) Regular Grammar
Regular grammars possess the most restrictive production rules within the Chomsky Hierarchy. These rules are confined to specific formats, such as a non-terminal deriving a terminal followed by a non-terminal or just a terminal. This high level of restriction limits their expressive power to regular languages, which can be recognized by machines with the simplest form of memory, finite automata.

34. (C) No algorithm can solve it for all inputs in finite time.
A problem is termed unsolvable or undecidable if it is mathematically proven that no single algorithm can exist that provides a correct yes or no answer for every possible input in finite time. This concept signifies a fundamental barrier in computation, showing that some questions cannot be resolved by any programmed machine, regardless of the complexity of the algorithm used.

35. (C) P-2, Q-1, R-1, S-2
Decidability varies across different computational problems. The halting problem for Turing machines and the Post Correspondence Problem are both famously undecidable, meaning no general algorithm can solve them. Conversely, the membership problem for context-free languages and the emptiness problem for finite automata are decidable, as there are established algorithms that can provide a correct result for any given instance.

36. (D) Match symbols by replacing them and moving head.
To recognize a language with balanced counts of two symbols using a Turing machine, a effective strategy involves a systematic matching process. The machine can replace one instance of the first

Downloaded from www.pdf.tube

ambiguities that a deterministic machine cannot handle with a single stack.

19. (C) A → BC or A → a
Chomsky Normal Form is a simplified way of representing context-free grammars where every production rule follows a strict format. Rules must either derive exactly two non-terminals or a single terminal symbol. This standardization is particularly useful for theoretical proofs and the implementation of efficient parsing algorithms, such as the CYK algorithm, which determines if a string belongs to a language.

20. (A) P-3, Q-4, R-1, S-2
Different classes of automata are designed to recognize specific levels of the language hierarchy. Finite automata recognize regular languages, while pushdown automata use stacks for context-free languages. Linear bounded automata handle context-sensitive languages by restricting tape usage to the input size. Finally, Turing machines are capable of recognizing recursively enumerable languages, the most general class in the hierarchy.

21. (C) A → aα, where 'a' is a terminal and α is a non-terminal string.
Greibach Normal Form is a specific restriction on context-free grammar productions where every rule begins with a terminal symbol followed by a potentially empty string of non-terminals. This form ensures that each step in a derivation consumes exactly one input symbol, making it highly compatible with pushdown automata and certain types of top-down parsing strategies for language recognition.

22. (B) A string in the language has multiple parse trees.
An ambiguous grammar is one that can generate more than one distinct parse tree for at least one string within the language it describes. This occurs when the production rules allow multiple paths of derivation, potentially leading to different structural interpretations of the same code. Resolving such ambiguity is critical in compiler design to ensure consistent and predictable program execution.

23. (D) A parse tree.
To visualize the structural derivation of an input string from the start symbol of a grammar, a parse tree is the most effective tool. It provides a hierarchical representation showing how production rules are applied to generate tokens. This graphical format helps developers understand operator precedence and associativity during the parsing process, facilitating the debugging of complex syntactic structures.

24. (C) (i), (iv), and (v) only
Context-free languages are closed under several fundamental operations, including union, concatenation, and the Kleene star. This means that applying these operations to context-free languages always results in another context-free language. However, unlike regular languages, they are notably not closed under intersection or complementation, as these operations can produce languages that require more power than a pushdown automaton provides.

25. (C) A finite control unit.
A standard Turing Machine is a theoretical model of computation consisting of an infinite tape divided into cells, a read/write head that moves along the tape, and a finite control unit. The finite control acts like a central processing unit, determining the machine's actions based on its current state and the symbol being read from the tape cell.

26. (A) Finite Automaton.
Many variations of the Turing Machine exist, such as those with multiple tapes, non-deterministic transitions, or tapes that are infinite in both directions. Despite these structural differences, all these models are equivalent in computational power to the standard single-tape machine. In contrast, a finite automaton is a strictly less powerful model because it lacks any external memory mechanism.

27. (A) Both Statement I and Statement II are true.
There is a complete equivalence between context-free grammars and pushdown automata in terms of the languages they can describe and recognize. Every context-free grammar has a corresponding pushdown automaton that accepts the same language, and conversely, any language recognized by a pushdown automaton

systematically processing tokens according to the language's syntax rules.

46. (A) Left-to-right scan, Leftmost derivation, 1 symbol lookahead.
In compiler design, LL(1) refers to a specific class of parsers and grammars. The first L stands for scanning the input from left to right, the second L indicates the generation of a leftmost derivation, and the 1 signifies the use of a single lookahead symbol to make parsing decisions. This combination allows for efficient, predictable parsing without backtracking.

47. (A) P, S, Q, R
Bottom-up parsers can be ordered by their ability to handle different grammar complexities. LR(0) is the simplest and least powerful, followed by SLR(1), which uses follow sets to resolve conflicts. LALR(1) offers a balance by merging states for efficiency, and LR(1) is the most powerful, using explicit lookaheads for each item to handle the widest range of context-free grammars.

48. (A) Uses fewer states in its table.
The primary advantage of LALR(1) over LR(1) parsing is the significant reduction in the size of the parsing table. LALR(1) achieves this by merging states in the LR(1) automaton that share the same core sets of items. While this occasionally introduces conflicts, it results in a much more memory-efficient parser that still retains most of the power of the LR(1) method.

49. (B) Shift-reduce parsing.
Bottom-up parsing is frequently referred to as shift-reduce parsing. This name describes the two fundamental actions the parser performs: shifting input tokens onto a stack and reducing sequences of symbols on the stack to a corresponding non-terminal based on grammar rules. This process continues until the entire input is reduced to the start symbol, confirming that the string is syntactically correct.

50. (B) P-4, Q-2, R-1, S-3
Different parsing strategies are characterized by their approach to tree construction and table management. Top-down parsing begins at the start symbol, while bottom-up parsing reduces the input toward the start symbol. LL(1) is a common predictive top-down method using lookahead. LALR(1) is a bottom-up technique that merges states from the more complex LR(1) method to optimize memory usage in tables.

51. (C) To verify static logic and types.
Semantic analysis is the phase of compilation where the compiler goes beyond syntax to ensure that the program's logic is sound. It verifies that variables are used consistently with their types and that operators are applied to compatible operands. This phase is crucial for detecting errors that are syntactically correct but logically impossible or undefined within the rules of the language.

52. (A) An attribute whose value is computed from its children.
In the context of an attribute grammar, a synthesized attribute is one whose value at a specific parse tree node is determined by the attribute values of its children. This information flows upward from the leaves of the tree toward the root. Such attributes are commonly used to calculate expression values or the combined type information of nested structural components.

53. (A) Both A and R are true and R is the correct explanation of A.
The Halting Problem's undecidability is proven through a classic proof by contradiction. If a machine could decide whether any program halts, one could design a paradoxical machine that behaves contrary to the decider's prediction. Specifically, the new machine would loop if the decider said it would halt and vice versa. This logical impossibility confirms that no such universal decider exists.

54. (A) Parent and/or siblings.
An inherited attribute is a type of attribute in a syntax-directed definition whose value at a parse tree node is calculated from the attributes of its parent or its siblings. This represents a downward or lateral flow of information within the tree structure. Inherited attributes are useful for passing context-specific information, such as variable types or scope levels, to various nested sub-components.

55. (B) An S-attributed definition.
A syntax-directed definition is categorized as S-attributed if it exclusively utilizes synthesized

Click Here to Buy book on Amazon

attribute in the cycle can be calculated without first knowing the value of another, the entire set becomes uncomputable. This is a critical error in syntax-directed definitions, as it prevents the compiler from establishing a valid evaluation sequence.

113. (C) Step consumes input.
The primary advantage of converting a grammar to Greibach Normal Form is that it ensures every step of a derivation consumes exactly one symbol of the input string. This direct correspondence between grammar rules and input tokens simplifies the design of pushdown automata and certain types of parsers. It guarantees that the parser makes steady progress through the input, facilitating efficient recognition.

114. (A) P-3, Q-1, S-2, T-4
Each phase of a compiler produces a specific output that serves as input for the next stage. The lexical analyzer outputs a stream of tokens, which the syntax analyzer uses to build a parse tree. The semantic analyzer then adds logical information to produce an annotated parse tree. Finally, the intermediate code generator translates this structure into a machine-independent format.

115. (B) Left recursive.
A grammar is not suitable for LL(1) parsing if it contains left-recursive production rules. In this top-down approach, a left-recursive rule would cause the parser to enter an infinite loop, attempting to expand the same non-terminal indefinitely without ever consuming an input token. To use such a grammar with an LL(1) parser, it must first be transformed to eliminate left-recursion.

116. (C) Code Optimization
The code optimization phase is the part of a compiler most directly responsible for the efficient assignment of variables to a limited set of hardware registers. This process, often called register allocation, uses techniques like graph coloring to ensure frequently used variables are kept in high-speed registers while minimizing memory accesses. This optimization is critical for producing high-performance machine code.

117. (D) Statement I is false, but Statement II is true.
Languages requiring a balance between two groups of symbols, such as an equal number of two characters, are context-free because they can be handled by a stack. However, when the balance must extend to three groups, the language is no longer context-free, as a single stack cannot track multiple independent counts simultaneously. Simple balanced languages are context-free, but three-way balances.

118. (D) Symbol Table
Intermediate representations like syntax trees, postfix notation, and three-address code are used by compilers to provide a simplified version of the source program for analysis. In contrast, the symbol table is a central data structure used to store information about identifiers across all phases. While it is essential for the compilation process, it is not considered a form of intermediate code.

119. (A) Globals cover procedures.
The key difference between global and local optimization is the scope of the analysis. Local optimization is restricted to a single basic block of code where the flow of control is linear. Global optimization, however, analyzes the entire control-flow graph of a procedure. This broader scope allows for powerful improvements, such as moving computations out of loops that span multiple blocks.

120. (C) Halt within 100 steps.
While most general questions about the long-term behavior of Turing machines are undecidable, specific questions about limited execution are decidable. For instance, determining if a machine halts within a fixed number of steps, like 100, is a simple task that can be solved by simulating the machine for that exact duration. This provides a definitive answer, unlike the general halting problem.

121. (B) Dynamic call nesting.
An activation tree is a dynamic representation of a program's execution, where each node corresponds to a specific invocation of a procedure. The hierarchical structure of the tree shows exactly which procedure called another and in what order these calls occurred during the runtime. This model captures the nested nature

as it avoids the overhead of copying the entire argument value.

65. (A) Track identifiers and scopes.
The symbol table is a central data structure used by a compiler to keep track of all identifiers encountered in the source program. It stores essential information about each name, such as its data type, its location in memory, and its scope within the program's block structure. This centralized repository is accessed throughout various phases of compilation to perform type checking.

66. (A) Three-address code.
Three-address code is a popular machine-independent intermediate representation used in many compilers. It breaks down complex source code expressions into a sequence of simple instructions, each involving at most three operands. This format provides a high degree of abstraction, making it easier to perform various code optimizations before the final translation into the machine code specific to the target processor.

67. (C) Entering data in symbol table.
When an intermediate code generator processes variable declarations, its primary task is to update the symbol table with information about the new identifiers. This involves recording the name, type, and size of each variable, which will be used later during semantic checks and memory allocation. This phase ensures that subsequent references to these variables can be correctly mapped and validated.

68. (A) t1 = c * d; t2 = b + t1; a = t2;
A complex assignment statement is translated into three-address code by breaking it down into smaller, discrete steps. Temporary variables are often introduced to store the results of intermediate calculations, such as the product of two variables. These steps must follow the correct operator precedence and associativity to ensure that the final result stored in the target variable matches logic.

69. (B) Jump if A is false.
In short-circuiting boolean evaluation, the compiler generates intermediate code that stops checking conditions as soon as the final result is determined. For a logical and expression, if the first part is false, the entire expression must be false. The generated code will include a conditional jump that skips the evaluation of the second part, improving performance and avoiding unnecessary logic.

70. (D) Re-parsing the body.
Generating intermediate code for a procedure call involves several steps to manage control flow and data. This includes pushing arguments onto the stack, saving the return address, and jumping to the procedure's entry point. However, re-parsing the body of the procedure is not part of this process, as the body has already been parsed and translated into intermediate code earlier.

71. (A) Nodes are blocks and edges are flow.
A control-flow graph is a graphical representation used in code optimization to model all possible paths of execution through a program. The nodes of the graph represent basic blocks, which are sequences of instructions without jumps. The directed edges represent the flow of control between these blocks, allowing the compiler to analyze and improve efficiency by identifying unreachable code.

72. (A) Collect possible program values.
Data-flow analysis is a fundamental technique used in compiler optimization to gather information about how values are computed and used throughout a program. By tracking the definitions and uses of variables across the control-flow graph, the compiler can identify opportunities for improvements, such as removing redundant calculations, thereby enhancing the overall performance of the generated code for the target.

73. (A) Local optimization.
Local optimization is a category of code improvement that focuses on a small, continuous sequence of instructions known as a basic block. Because there are no jumps into or out of the middle of the block, the compiler can safely perform optimizations like constant folding or redundant instruction removal. These changes are isolated and do not require analysis of global flow.

74. (B) Replicates the loop body.

Downloaded from www.pdf.tube

attributes. In this model, every attribute value is derived solely from the node's children, allowing for a straightforward evaluation process during a single bottom-up pass of the parse tree. This simplicity makes S-attributed definitions highly compatible with shift-reduce parsers, as attributes can be computed upon each reduction.

56. (B) (i), (ii) and (iii) only
Syntax-directed definitions provide a formal framework for associating semantic rules with the productions of a context-free grammar. They are instrumental in specifying both the structure and the meaning of programming languages. While many can be evaluated efficiently, not all SDDs can be processed in a single pass of a bottom-up parser, as some may involve complex dependencies between attributes.

57. (D) Order attribute evaluation.
A dependency graph in semantic analysis is a directed graph that illustrates the relationships between various attributes in a parse tree. Each edge indicates that one attribute's value must be calculated before another. This graph is essential for determining a valid evaluation order, ensuring that every attribute is computed using the necessary data from its predecessors, thus preventing circularity.

58. (C) Synthesized attributes and attributes inherited from the parent or from siblings to its left.
An L-attributed definition is a class of syntax-directed definitions that allows for both synthesized attributes and a restricted type of inherited attributes. Specifically, it permits an attribute to depend on values from the parent node or from siblings located to its left in the parse tree. This restriction ensures that attributes can be evaluated during a single left-to-right, top-down traversal.

59. (B) Verify operand compatibility.
The primary responsibility of the type-checking process within the semantic analysis phase is to ensure that all operations are performed on compatible data types. By verifying operand compatibility, the compiler can prevent many common programming errors, such as attempting to add a string to an integer. This step ensures that the resulting intermediate code will be logically consistent.

60. (D) Parse tree
Run-time storage organization for block-structured languages typically includes segments for executable code, a stack for procedure calls, and a heap for dynamic memory. A parse tree, however, is an intermediate structure used by the compiler during the translation process. Once the compiler has generated the target code, the parse tree is no longer needed and does not exist in memory.

61. (A) Dynamic function call sequences.
An activation tree is a graphical representation used to visualize the sequence and nesting of procedure calls during a program's execution. Each node in the tree represents an individual activation of a procedure, and the edges show which procedure called another. This model captures the dynamic flow of control and helps illustrate how activation records are pushed onto and popped off the stack.

62. (C) Global symbol table.
Activation records are designed to store information specific to a single invocation of a procedure, such as local variables, the return address, and passed parameters. They facilitate recursion and nested calls by isolating the data for each call. In contrast, the global symbol table contains information about identifiers across the entire program and is typically managed outside of individual activation records.

63. (A) Recursion.
The use of a stack for managing activation records is essential for supporting languages that allow recursion. Each time a procedure calls itself, a new activation record is pushed onto the stack to hold its unique local data and return address. This structure ensures that when a recursive call completes, the program can correctly restore the previous state and continue execution.

64. (A) Pass-by-reference.
Pass-by-reference is a parameter-passing mechanism where the memory address of the argument is passed to the called procedure. This creates an alias for the original variable, allowing any modifications made within the procedure to directly affect the variable in the calling scope. This method is efficient for large data structures

NODIA

simulate the behavior of any other machine. While they are highly powerful, they cannot solve the undecidable halting problem, and non-deterministic versions are equivalent to deterministic ones.

84. (B) Pushing symbols onto stack.
In bottom-up parsing, specifically LR parsing, a shift action occurs when the parser consumes the next input token and pushes it onto its internal stack. This action represents progress through the input string as the parser builds up a sequence of symbols that it will eventually reduce to a non-terminal. Shifting is the primary way the parser moves forward during processing.

85. (C) Multiplication of y, z.
Intermediate code generators must respect the standard operator precedence of the source language when translating expressions. In an expression involving both addition and multiplication, the multiplication must be performed first. The compiler will generate instructions to calculate the product of the relevant operands before adding that result to the other term, ensuring that the final computation follows established mathematical rules.

86. (A) P-3, Q-4, R-1, S-2
Various code optimization techniques serve specific purposes in improving program performance. Loop-invariant code motion moves constant calculations out of loops, while constant folding evaluates expressions with constant values at compile time. Strength reduction replaces expensive operations with cheaper ones, such as changing multiplication to addition. Dead code elimination removes instructions that are never executed or whose results are never used.

87. (C) Rules depend on context symbols.
A context-sensitive grammar is defined by production rules where the transformation of a non-terminal depends on the symbols surrounding it. This means that a symbol can only be replaced if it appears in a specific context. Mathematically, the length of the string on the left side of a production must be less than or equal to the length on the right.

88. (A) Polynomial time.
In computational complexity, a problem is considered tractable if it can be solved by an algorithm with a running time that is a polynomial function of the input size. This classification suggests that the problem is practically solvable for large inputs, even if the degree of the polynomial is high. Most common algorithms in computer science fall into this category.

89. (D) Left recursive.
A production is called left recursive when the nonterminal on the left side appears as the leftmost symbol on the right side of the rule. In expr → expr - expr, the symbol expr immediately repeats on the left of the production's right-hand side. This creates left recursion, which can cause issues for top-down parsers and often needs to be eliminated during grammar design.

90. (A) It may introduce reduce conflicts.
LALR(1) parsers are created by merging states from an LR(1) parser that share the same core items, differing only in their lookahead symbols. While this reduces the table size, a potential downside is that the merging process can introduce reduce-reduce conflicts. This happens when the combined lookahead sets make it impossible for the parser to distinguish between two different reductions.

91. (C) Semantic Analysis
The semantic analysis phase of a compiler is responsible for ensuring that the source program follows logical rules that go beyond simple syntax. One of its key tasks is to verify that every variable is declared before it is used. This requires checking identifiers against the symbol table to ensure they have been properly defined within the current scope, preventing runtime failures.

92. (A) Separation of front and back end.
Using an intermediate code representation allows for a clean separation between the compiler's front end, which handles language-specific analysis, and the back end, which handles machine-specific code generation. This modularity makes it much easier to retarget a compiler for different processors or to support multiple programming languages by reusing the same optimization and code generation components, reducing development effort.

Click Here to Buy book on Amazon

Loop unrolling is an optimization strategy that replicates the body of a loop multiple times to reduce the overhead of the loop control logic. By decreasing the frequency of jump instructions and test conditions, the compiler can often improve execution speed. Additionally, unrolling can create more opportunities for other optimizations, such as better instruction scheduling and efficient pipeline use.

75. (A) Replacing instruction sequences.
Peephole optimization is a local code improvement technique that examines a small, sliding window of generated instructions. The compiler looks for specific patterns or sequences that can be replaced with more efficient equivalents, such as removing redundant moves or replacing slow instructions with faster ones. This process is often applied repeatedly to the final machine code to squeeze out extra performance.

76. (B) Minimize pipeline stall issues.
Instruction scheduling is a crucial phase in back-end compilation that focuses on reordering the final machine instructions to maximize efficiency. Its primary goal is to minimize pipeline stalls caused by data dependencies or resource conflicts within the processor. By carefully arranging the sequence of operations, the compiler ensures that hardware components are kept busy, leading to faster execution times on modern architectures.

77. (C) Lexical Analysis - Semantic logic.
Lexical analysis is the initial phase where code is divided into tokens, whereas semantic analysis is responsible for checking logical correctness and type consistency. Associating lexical analysis with semantic logic is incorrect, as these phases serve different purposes. The Pumping Lemma is for proving non-regularity, Chomsky Normal Form aids in parsing algorithms, and the symbol table is the repository for identifier data.

78. (D) Statement I is false, but Statement II is true.
Deterministic finite automata are strictly defined so that each state has exactly one transition for every symbol in the alphabet. In contrast, non-deterministic finite automata have more flexibility, allowing for multiple possible transitions from a single state on the same input symbol. Furthermore, NDFAs can include epsilon transitions, which allow the machine to move between states without consuming any characters.

79. (B) Context-Free only
A language consisting of strings with an equal number of two different symbols, like 0s and 1s, is not regular because it requires counting to ensure the balance. However, it can be recognized by a pushdown automaton using a stack to track the excess of one symbol over the other. Therefore, this language is context-free, but remains beyond the power of automata.

80. (A) Both A and R are true and R is the correct explanation of A.
S-attributed definitions are well-suited for bottom-up parsing because they only use synthesized attributes. As the parser reduces a sequence of symbols to a non-terminal, it has already completely processed the corresponding subtrees. This means the values for the children's attributes are available and can be immediately used to compute the attribute for the parent, allowing for an efficient process.

81. (B) All strings over {a, b}.
A grammar with productions that allow a start symbol to recursively generate either 'a' or 'b' followed by itself, or to produce an empty string, can create any sequence of those characters. This structure encompasses all possible combinations, including the empty string. Thus, the language generated is the set of all strings over the given alphabet, representing universal regular strings.

82. (A) 2^N states.
The subset construction algorithm is used to convert any non-deterministic finite automaton into an equivalent deterministic one. While the two machines recognize the same language, the number of states in the resulting DFA can grow exponentially relative to the number of states in the original NFA. In the worst-case scenario, an NFA with N states could lead to a very large DFA.

83. (B) (iii) and (iv) only
The Church-Turing thesis establishes that Turing machines are a universal model of computation, capable of solving any problem that has an algorithmic solution. Furthermore, a Universal Turing Machine is a specific model that can

NODIA

with the corresponding non-terminal. This process continues until input is reduced to start.

103. (A) Both A and R are true and R is the correct explanation of A.
Left-recursion in a grammar is a significant issue for top-down parsing methods like recursive descent. This is because a left-recursive rule will cause the parser to repeatedly call the same function without ever consuming a token from the input. This leads to an infinite loop of recursive calls, eventually exhausting the system's memory or stack space and preventing successful processing.

104. (C) Dynamic allocation.
Memory allocation for objects whose lifetime is determined at runtime, such as those created with commands like malloc or new, is known as dynamic allocation. These objects are stored in a specialized region of run-time memory called the heap. Unlike stack-allocated data, heap-allocated memory must be explicitly managed or reclaimed by a garbage collector to prevent potential memory leaks.

105. (B) With bottom-up parsing.
S-attributed syntax-directed definitions are most effectively used in conjunction with bottom-up parsing. Since these definitions rely exclusively on synthesized attributes, the necessary information from children nodes is always available when a reduction occurs. This allows the compiler to perform all semantic calculations on the fly as it constructs the parse tree from the leaves upward, making it an efficient choice.

106. (A) Pumping lemma for CFLs.
To rigorously prove that a language like the one consisting of doubled strings is not context-free, mathematicians use the Pumping Lemma for context-free languages. This technique shows that long strings in such a language cannot be broken into sections and repeated without eventually creating a string that no longer fits the doubled format. This contradiction confirms language power limits.

107. (B) (i), (ii), (iv), and (v) only
Parsing techniques are broadly categorized into top-down methods like LL and recursive descent, and bottom-up methods like LR and LALR. LR parsers are generally more powerful and can handle a wider range of grammars than LL parsers. LALR is a specific version of LR that optimizes memory by merging states. However, recursive descent is a top-down parsing method.

108. (A) The equivalence problem.
The equivalence problem, which asks if two machines recognize the same language, has a different status depending on the model. For finite automata, this problem is decidable, as there are algorithms to compare their state transitions. However, for Turing machines, this problem is undecidable, meaning no general algorithm can ever be created to determine if two arbitrary programs are identical.

109. (A) Continuous code without jumps.
In the context of intermediate code and compiler optimization, a basic block is a continuous sequence of instructions with exactly one entry point and one exit point. This means there are no jumps into the middle of the block and no branch instructions within it, except possibly at the end. This clear structure makes basic blocks the unit for analysis.

110. (B) Contain 'aa'.
A regular expression that features a specific substring bracketed by universal character matches represents the set of all strings containing that substring. In this case, the sequence of two 'a' characters is surrounded by expressions that allow any combination of symbols of any length. Consequently, any string over the alphabet that includes the 'aa' pattern will be matched by expression.

NODIA

111. (A) Also recursive.
The class of recursive languages is closed under complementation. This means that if a language can be decided by a Turing machine that always halts, its complement can also be decided by a similar machine. The process involves simply swapping the accept and reject results of the original machine. This property ensures that the decidability of a language is mirrored.

112. (B) Attributes are uncomputable.
A cycle in an attribute dependency graph indicates that a group of attributes depends on each other in a circular fashion. Because no

93. (C) May be used before redefine.
A variable is considered live at a specific point in a program if its current value is needed for a calculation that occurs later in at least one possible execution path. Tracking live variables is essential for many code optimizations, particularly for register allocation. If a variable is no longer live, the memory or register it occupies can be safely reused.

94. (B) Recognizing lexemes.
In the lexical analysis phase, a finite automaton is used as a systematic engine for recognizing lexemes from the input character stream. The automaton transitions between states based on each character read, following rules defined by regular expressions. When the machine reaches an accepting state, it identifies a valid lexeme and outputs the corresponding token, providing a structured stream for the parser.

95. (B) Context-sensitive grammars.
Within the Chomsky Hierarchy, context-sensitive grammars correspond to the most powerful class of languages that are guaranteed to be recognized by a halting algorithm. These languages are a subset of the recursive languages, meaning that for any context-sensitive language, a Turing machine can be constructed that will always halt and correctly decide if a given string belongs to the language.

96. (A) (i), (ii) and (iv) only
A pushdown automaton is essentially a finite state machine that uses an external stack to provide additional memory. This structure allows it to recognize the entire class of context-free languages. Each transition in a PDA is determined by the current state, the input symbol being read, and the symbol currently on top of the stack, enabling nested structures.

97. (B) Function gets a copy.
Pass-by-value is a common parameter-passing mechanism where the called function receives an independent copy of the argument's data. Because the function operates on this copy, any changes made to the parameter within the function's scope do not affect the original variable in the calling code. This approach ensures data isolation and prevents unintended side effects during program execution.

98. (C) Instruction scheduling.
Instruction scheduling is a prime example of machine-dependent code optimization. It involves reordering the final machine instructions to match the specific pipeline and resource characteristics of a particular processor. Unlike machine-independent optimizations like constant folding or loop-invariant code motion, instruction scheduling must be tailored to the target hardware to effectively minimize pipeline stalls and maximize the execution throughput.

99. (B) Terminal or ε.
In a parse tree, which graphically represents the derivation of a string from a grammar, the leaf nodes correspond to the symbols that actually appear in the final input string. These are the terminal symbols of the grammar, representing the tokens identified during lexical analysis. Additionally, a leaf node can represent the empty string symbol, signifying reduction to nothing.

100. (A) R, S, P, Q
A typical compiler processes source code through a specific sequence of phases to ensure correctness and efficiency. It begins with lexical analysis to identify tokens, followed by syntax analysis to build a structural representation like a parse tree. Next, semantic analysis checks for logical and type errors. Finally, the code generation phase translates the validated program structure into the target machine code.

101. (B) Not solvable in polynomial.
An intractable problem is one that, while solvable in principle, cannot be solved by any known algorithm in polynomial time. These problems typically require exponential or super-polynomial time as the input size grows, making them practically impossible to solve for all but smallest instances. This category is distinct from undecidable problems, which cannot be solved by any algorithm regardless of time.

102. (A) Substring ready for reduction.
In bottom-up or shift-reduce parsing, a handle is a specific substring on top of the parser's stack that matches the right-hand side of a grammar production rule. Identifying the correct handle is crucial because it indicates that a reduction can be performed, replacing the sequence of symbols

Downloaded from www.pdf.tube

of procedure activations, making it a valuable debugging tool.

122. (D) Enumerable
In the Chomsky hierarchy, unrestricted grammars are the most powerful class, corresponding to recursively enumerable languages. These grammars have no limitations on their production rules, allowing for any transformation of symbol strings. Consequently, they can generate any language that can be recognized by a standard Turing machine, representing the absolute limit of what is computationally possible using algorithmic formal processes.

123. (B) Handle
During bottom-up parsing, a handle is a specific substring found on top of the parser's stack that matches the right-hand side of a grammar rule. Recognizing a handle signals that the parser can perform a reduction, replacing that sequence of symbols with the non-terminal on the left side of the rule. This is a key step in building derivations.

124. (A) R, S, Q, P
In a typical program, a variable follows a specific lifecycle. It is first declared, providing the compiler with its type and scope. Memory is then allocated to hold its data when its containing procedure is activated. The variable is used throughout the computation as needed. Finally, when the variable's lifetime ends, the memory it occupied is deallocated to be reused.

125. (C) Semantic Analysis
The task of resolving an overloaded operator, such as determining if a plus sign represents integer or floating-point addition, is a core responsibility of the semantic analysis phase. By checking the types of the operands involved, the compiler can assign the correct meaning to the operation. This ensures that subsequent code generation uses the appropriate machine instructions for the types.

126. (C) It is not enumerable.
If a language is recursively enumerable but not recursive, its complement is guaranteed to be non-enumerable. This is a fundamental result in computation theory. If both a language and its complement were recursively enumerable, a machine could simply run two simulations in parallel to decide membership, which would make the language recursive. The absence of decidability forces this non-enumerable status.

127. (D) LR(1) - Bottom-up
Parsing techniques are classified by how they build the parse tree. LR(1) is a sophisticated bottom-up parsing method that uses a stack and lookahead to reduce input tokens toward the start symbol. Other pairings are often incorrect; for example, recursive descent is a top-down method, not bottom-up, and shift-reduce is also a bottom-up category. Identifying these is essential.

128. (C) Strength reduction
Strength reduction is an optimization technique that replaces a computationally expensive operation with a simpler and faster one that produces the same result. A classic example is replacing a multiplication operation, which may take several clock cycles, with an addition operation that can be executed much more quickly. This improvement is particularly effective when applied within loops executed many times.

129. (A) Both A and R are true and R is the correct explanation of A.
While context-free languages are closed under union and concatenation, they are notably not closed under intersection. This is demonstrated by the fact that two context-free languages, each tracking a different pair of symbols, can have an intersection that requires a three-way balance. Since such a balance is beyond the power of a single stack, the resulting language is context-sensitive.

130. (A) Reach non-local data.
An access link, also known as a static link, is an important component in the activation record of languages that support nested procedures and lexical scoping. Its primary purpose is to provide a way for a procedure to access non-local variables located in its outer, enclosing scopes. By following these links, the program can traverse the stack for identifiers.

131. (B) CFLs close on intersection.
The class of context-free languages is not closed under intersection, meaning that the intersection of two such languages is not always context-free. This is a common misconception, as regular

NODIA

languages do possess this closure property. Other statements are correct: regular languages are a subset of context-free ones, they are closed under complement, and the union of context-free languages is context-free.

132. (B) Ambiguous
A grammar is classified as ambiguous if there is at least one sentence in the language it defines that can be derived in more than one way, leading to multiple distinct parse trees. This ambiguity can cause uncertainty in structural interpretation, potentially leading to different behaviors. Compilers resolve this by using precedence and associativity rules to ensure a single consistent interpretation.

133. (D) All are correct.
Compilers and interpreters are different approaches to program execution. A compiler translates the whole program into machine code beforehand, resulting in faster runtime performance. An interpreter processes instructions one by one, offering flexibility for debugging. Despite these differences, both methods rely on fundamental front-end processes, including lexical and syntax analysis, to understand and validate the source code's structure and meaning.

134. (C) Turing Machine
The Turing Machine is the most powerful computational model in the hierarchy of automata. It is capable of simulating any algorithmic process and can recognize the broadest class of languages, known as recursively enumerable languages. Other models like finite automata and pushdown automata are more restricted because they have limited memory structures, making them unable to solve many complex problems.

135. (A) Postfix translation.
In a syntax-directed translation scheme, if all semantic rules are placed exclusively at the end of the right-hand side of each production, it is known as a postfix translation. This specific arrangement is efficient for bottom-up parsing. It allows the compiler to execute semantic actions immediately after a reduction occurs, as all necessary data from children symbols is already available.

136. (C) Semantic Analysis
Semantic analysis is the phase of compilation specifically tasked with verifying that the program's logic and data usage follow the language's rules. This includes detecting errors like type mismatches, where a value of one type is incorrectly assigned to a variable of a different type. While syntax analysis ensures code structure is correct, semantic analysis ensures operations make sense.

137. (A) One op, three addresses.
A standard three-address instruction is designed to be simple and regular, containing at most one operator and three addresses. These addresses typically represent two operands for the operation and one destination where the result is stored. This predictable format simplifies code generation and makes optimization techniques more straightforward to implement, as each instruction represents a single, discrete step in computation.

138. (A) Not in that class.
The Pumping Lemma defines a necessary property that every language in a certain class, such as regular or context-free, must possess. If it is mathematically proven that a specific language does not satisfy the conditions of the Pumping Lemma, the only logical conclusion is that the language does not belong to that class. This makes it a proof tool.

139. (C) P-3, Q-1, R-4, S-2
Various techniques are used throughout compilation to handle grammar transformations, state machines, and code generation. Backpatching is used to resolve jumps in control flow, while left factoring prepares grammars for predictive parsing. Peephole optimization improves local code efficiency, and subset construction is the standard method for converting non-deterministic finite automata into deterministic ones. These diverse methods ensure correctness.

140. (B) Ambiguous CFL.
While every regular language can be accepted by both deterministic and non-deterministic finite automata, the same is not true for context-free languages and pushdown automata. There exists a class of context-free languages that can only

Click Here to Buy book on Amazon

be recognized by non-deterministic pushdown automata. These languages are often associated with inherent ambiguity that cannot be resolved by a deterministic machine using a stack.

141. (B) Heap manager.
The heap manager is a critical part of the run-time system responsible for allocating and deallocating memory for objects with dynamic lifetimes. Unlike stack-based data, which follows strict procedure call order, heap-allocated objects can be created and destroyed in any order. The manager must track available memory segments to fulfill requests and ensure that memory is efficiently used overall.

142. (A) Shift-reduce conflict.
A shift-reduce conflict is a common issue in LR parsing where the parser encounters a state and a lookahead symbol that allow for two different actions. Specifically, it could either shift the symbol onto the stack or reduce current stack contents using a grammar rule. This ambiguity must be resolved, often by using precedence rules, to ensure a single path.

143. (A) (i), (iii) and (iv) only
Code optimization aims to improve program efficiency while strictly maintaining its original meaning. While many optimizations reduce code size, some techniques, such as loop unrolling, can actually increase it to achieve faster execution speeds. Loop-invariant code motion is a powerful global optimization. In contrast, peephole optimization is a local technique, not global, as it only examines small windows.

144. (A) It simulates any TM.
The fundamental concept of a Universal Turing Machine is its ability to simulate the behavior of any other Turing Machine. By taking an encoded description of another machine and an input string, the UTM can replicate the exact computational steps that the original machine would have taken. This demonstrates that a single machine is capable of performing any algorithmic computation.

145. (A) Text is substituted.
Call-by-name is a parameter-passing mechanism where the actual text of the argument is substituted for the formal parameter in the procedure's body. The argument expression is then re-evaluated every time the parameter is accessed during execution. This approach can lead to complex and inefficient behavior if the expression is complicated or involves side effects, as it lacks pre-evaluation.

NODIA

146. (A) Both Statement I and Statement II are true.
LR(1) parsers are inherently more powerful than LALR(1) parsers because they maintain distinct states for every unique combination of items and lookahead symbols. LALR(1) simplifies this by merging states with the same core items, which can occasionally introduce reduce-reduce conflicts not present in the original model. This trade-off between parsing power and table size is a central consideration.

147. (A) It handles common parts.
A Directed Acyclic Graph is often used as a more efficient alternative to a syntax tree because it can identify and merge common subexpressions. In a DAG, multiple parents can point to the same child node, allowing a single representation for a value used in multiple parts of an expression. This compact structure simplifies various optimizations and reduces redundancy.

148. (B) Subset
Regular languages are highly robust, showing closure under union, intersection, and complementation. This means performing these operations on regular sets always results in another regular set. However, they are not closed under the subset operation. A subset of a regular language is not necessarily regular itself, as it could be any language, including one that is non-regular or undecidable.

149. (B) Undecidable
Determining if two arbitrary context-free grammars generate exactly the same language is a famously undecidable problem. There is no general algorithm that can compare any two grammars and always provide a correct yes or no answer in a finite amount of time. This limitation highlights the complexity of context-free languages compared to regular languages, where this problem is decidable.

150. (B) R, S, Q, P

A compiler systematically processes source code through several distinct stages. It first performs tokenization to identify basic language elements, which are then used to construct a structural parse tree. Once the structure is established, the compiler performs type checking to ensure logical consistency. Finally, after various optimizations, it performs register allocation to map program variables to physical machine registers.

151. (A) TMs compute everything.
The Church-Turing Thesis is a foundational hypothesis which proposes that any process that can be described as an algorithm can be performed by a Turing machine. It effectively defines the boundary of what is computable, suggesting that no machine can exceed the computational limits of this universal model. This thesis underpins much of theoretical computer science and our understanding of computation.

152. (A) Handle swap for non-terminal.
In bottom-up LR parsing, the reduce action is the process of recognizing a complete structural component of the grammar on the stack. When a sequence of symbols matches the right-hand side of a production rule, the parser replaces that entire sequence with the rule's corresponding non-terminal. This swapping of a handle for a higher-level symbol builds the grammar structure.

153. (A) Both A and R are true and R is the correct explanation of A.
The class of regular languages is a proper subset of context-free languages. Every regular language is automatically context-free because a pushdown automaton can simply ignore its stack and act like a finite automaton. Conversely, there are many context-free languages, such as the one requiring balanced groups of symbols, that have been proven beyond the power of automata.

154. (A) Accessing non-local data.
In certain run-time systems, a display is an auxiliary array of pointers used to facilitate efficient access to non-local variables in nested procedures. Each element in the display points to the activation record of the most recent instance of a procedure at a specific nesting depth. This allows the program to jump directly to the correct data environment without static links.

155. (A) Rules have no limits.
Unrestricted grammars are characterized by their complete lack of limitations on production rules. They allow any sequence of symbols on the left side of a rule to be replaced by any sequence on the right side. This absolute flexibility gives them the same computational power as a standard Turing machine, allowing them to generate the most general language class possible.

156. (A) Values stay constant.
Moving a computation out of a loop is a valid optimization only if the values of the operands involved remain constant throughout all iterations of the loop. If any variable used in the calculation is modified within the loop, the result could change, and moving it outside would alter the program's meaning. Ensuring this constancy is critical for maintaining logical correctness.

157. (C) RE and co-RE is Rec.
A fundamental result in formal language theory is that if both a language and its complement are recursively enumerable, then the language is guaranteed to be recursive. This is because we can simulate machines for both sets in parallel; one of them must eventually halt and provide an answer. This property links semi-decidability and full decidability within the hierarchy.

158. (A) Checking parenthesis.
The primary job of the lexical analyzer is to identify tokens and remove irrelevant elements like whitespace from source code. However, checking for balanced parentheses is a syntactic task, as it involves understanding nested structures and the order of symbols. This responsibility falls to the syntax analysis phase, which uses a parser and stack to ensure the code follows rules.

159. (A) Numerical value.
In a syntax-directed definition designed for a desk calculator, the most logical attribute to associate with an expression non-terminal is its numerical value. As the parser processes input, it can apply semantic rules to calculate and store this value for each sub-expression. This allows the final result for the entire expression to be synthesized from its lower-level components.

160. (B) Limited state count.

Downloaded from www.pdf.tube

The fundamental limitation of a finite automaton is that it possesses only a fixed and finite number of states, which restricts its memory capacity. This makes it impossible for the machine to recognize languages that require tracking an arbitrary count of symbols. Without the ability to count beyond its state limit, an automaton cannot recognize balanced, nested, or repeated structures.

161. (A) Intractable.
In computational complexity, a problem that is theoretically solvable but for which no polynomial-time algorithm has been found is classified as intractable. Such problems often have exponential time complexity, meaning that as input size increases, the time required to find a solution grows dramatically. While a computer can find the answer in principle, the practical time needed often exceeds reasonable limits.

162. (A) P-3, Q-4, R-1, S-2
The Chomsky Hierarchy pairs four levels of grammars with corresponding machine models. Unrestricted grammars are matched with Turing machines, which offer highest power. Context-sensitive grammars align with linear bounded automata, and context-free grammars correspond to pushdown automata using stacks. Finally, regular grammars, the most restrictive type, are recognized by finite automata, which have the simplest form of computational processing power.

163. (A) Elimination of left recursion.
The process of restructuring grammar rules to replace a production where a non-terminal calls itself at the start is called the elimination of left-recursion. By introducing a new non-terminal and an empty string option, the grammar is transformed into an equivalent right-recursive form. This modification is essential for making a grammar compatible with top-down parsing techniques, like recursive descent.

164. (A) Resuming caller code.
The return address is a critical piece of information stored in an activation record when a procedure is called. Its purpose is to record the exact location in the code where execution should resume once the current procedure completes its task. By storing this address, the program can correctly return control to the caller, ensuring the dynamic flow remains logical.

165. (D) Statement I is false, but Statement II is true.
Peephole optimization is a highly localized technique that operates by examining a small, fixed-size window of instructions, typically in final machine or assembly code. The compiler looks for specific patterns within this peephole that can be replaced with more efficient equivalents. Although simple, it is often applied repeatedly to achieve significant cumulative improvements in the speed and size of code.

166. (A) Choice between reductions.
A reduce-reduce conflict occurs in an LR parser when the machine reaches a state and sees a lookahead symbol that suggests two or more different grammar rules could be used for reduction. Because the parser can only choose one path, this ambiguity represents a conflict. This often indicates the underlying grammar is ambiguous or requires more lookahead symbols to be parsed.

167. (A) All S are also L.
In the categorization of syntax-directed definitions, every S-attributed definition is also an L-attributed definition. S-attributed definitions are a more restrictive class that only uses synthesized attributes, while L-attributed definitions allow for both synthesized and a restricted set of inherited attributes. This nested relationship means that the simpler S-attributed model is a special case within the broader L-attributed framework.

NODIA

168. (A) No universal decider exists.
The fact that the Halting Problem is undecidable means that it is fundamentally impossible to create a single, universal computer program that can correctly determine whether any other arbitrary program will eventually stop or run forever. This discovery establishes a hard limit on what can be achieved through computation, proving that some well-defined questions are forever beyond algorithmic solutions.

169. (A) (i), (ii), and (iv) only
In compiler design, the symbol table is a vital data structure that remains active throughout the entire process. Syntax analysis correctly outputs a

corresponding virtual machine. This approach provides a high level of portability, allowing the same compiled program to run on diverse hardware.

189. (A) Lexical scoping rules.
A static link, also known as an access link, is a pointer used in activation records to implement lexical scoping in languages with nested procedures. Its primary purpose is to point to the activation record of the procedure's static parent in the source code. By following these links, a procedure can correctly access non-local variables defined in outer scopes, adhering to rules.

190. (A) Syntax trees are cleaner.
While both trees represent the structure of a program, a syntax tree is a more streamlined version than a parse tree. A parse tree captures every detail of the grammar's derivation, including intermediate non-terminals that might not be necessary. A syntax tree, in contrast, focuses on the core logical structure, removing redundant information and making it easier for subsequent phases.

191. (A) R, S, Q, P
The power of computational models can be ordered from most capable to least capable. Turing machines are the most powerful, capable of recognizing any enumerable language. Non-deterministic pushdown automata are next, covering all context-free languages, while deterministic versions are slightly more restricted. Finally, finite automata are the least powerful, recognizing only regular languages. This accurately reflects expressive capacity.

192. (C) Static rule checks.
The primary responsibility of the semantic analyzer is to perform static logic and rule checks on the source program. After the syntax has been validated, the semantic analyzer ensures that the operations make sense, such as verifying that variables are used with correct and compatible types. This phase is essential for identifying errors that are logically impossible or prohibited.

193. (A) {0^n 1^n where n ≥ 0}
A grammar with the production S → 0S1 | ε creates strings where every '0' is followed by a matching '1' later in the sequence. Each recursive application of the rule adds another pair, and the ε rule allows the process to terminate. This results in a language of all strings that have a sequence of '0's followed by equal '1's.

194. (C) Statement I is true, but Statement II is false.
Code optimization is an optional phase in many compilers, included to improve the speed or memory efficiency of the final program. While it is highly beneficial, a compiler can produce a correct executable without it. However, it is a misconception that all optimizations reduce size; some techniques, like loop unrolling, can increase code size to achieve faster execution.

195. (B) Hash table stack.
For a block-structured language with nested scopes, a hash table stack is the most suitable data structure for managing the symbol table. This structure allows the compiler to easily push new scopes onto the stack and pop them off as it enters and leaves blocks. Within each scope, a hash table provides fast access to identifier information, effectively handling requirements.

196. (A) Simple.
In the context of LR parsing, SLR stands for Simple LR. The Simple designation refers to the way this parser uses Follow sets to resolve potential conflicts in its parsing table. While less powerful than the more sophisticated LR(1) or LALR(1) methods, an SLR parser is easier to construct and can handle a wide range of common programming language grammars.

197. (D) Complementation
Context-free languages are robust under several operations, including union, concatenation, and the Kleene star. This means performing these actions on context-free languages always results in another context-free language. However, they are notably not closed under complementation. The complement of a context-free language can be a more complex language that requires more expressive power than a pushdown automaton possesses.

198. (A) P-2, Q-4, R-1, S-3
Complexity classes categorize problems based on the resources required to solve them. Sorting an array is a classic example of a tractable problem

Click Here to Buy book on Amazon

179. (A) Both A and R are true and R is the correct explanation of A.
The class of context-free languages is closed under the union operation. This can be easily proven by a construction where a new grammar is created with a new start symbol that can derive the start symbols of two other grammars. Because this simple structural change results in a grammar generating the combined set of strings, union remains context-free.

180. (A) Minimize stall issues.
Instruction scheduling is a phase in back-end compilation with the primary goal of reordering machine instructions to minimize pipeline stalls. By carefully arranging the sequence of operations, the compiler can account for data dependencies and resource conflicts within the processor. This ensures that hardware components are used efficiently, leading to significant improvements in final program execution speed on modern architectures.

181. (A) Single symbol lookahead.
In the context of LR parsing, the 1 specifies that the parser uses exactly one lookahead symbol from the input stream to make its parsing decisions. This additional information allows the LR(1) parser to be more powerful and precise than the LR(0) model, which makes decisions based solely on stack state. The lookahead symbol helps resolve potential shift-reduce conflicts.

182. (A) No algorithm can exist.
A non-computational problem is one for which it has been mathematically proven that no possible algorithm can ever exist to solve it. These problems are not simply difficult or time-consuming; they are fundamentally beyond the reach of any programmable machine. Regardless of how much time or processing power is available, no systematic procedure can provide a correct answer for every instance.

183. (A) Downward type passing.
In semantic analysis, an inherited attribute is used to pass information downward from a parent node to its children or laterally between siblings. A classic example is the downward passing of type information in a declaration. The type determined at a higher level in the parse tree can be inherited by identifiers below it, ensuring correct data type associations.

184. (A) (b*ab*a)*b*
To represent a language with an even number of 'a's, a regular expression must ensure that every 'a' is paired with another 'a'. In between and around these pairs, any number of 'b's can occur. This structure ensures that no matter how many times the core pattern is repeated, the total count of 'a's will always be a multiple of two.

185. (A) Undecidable string match.
The Post Correspondence Problem is a well-known undecidable problem in formal language theory. It asks whether two lists of strings can be arranged in a specific sequence so that the concatenated result from the first list exactly matches the concatenated result from the second list. Its unsolvability demonstrates that even seemingly simple string matching tasks can be beyond any algorithmic solution.

186. (A) Merged core LR(1) states.
LALR(1) parsers are a sophisticated type of bottom-up parser generated by merging states from an LR(1) automaton that have identical core sets of items. This merging process results in a much smaller parsing table, making it more memory-efficient while still retaining significant power. While this can sometimes introduce reduce-reduce conflicts, LALR(1) remains the most common method in modern generators.

187. (B) (ii) and (iv) only
In the study of formal grammars, context-free grammars are more expressive than regular grammars. Unrestricted grammars are the most powerful, generating the entire class of recursively enumerable languages. However, some statements can be misleading; for example, empty string productions are handled differently in context-sensitive grammars, and Chomsky Normal Form specifically requires rules to derive either two non-terminals or a terminal.

188. (A) Virtual machine compiler.
A virtual machine compiler is a type of system that translates a source program into the instruction set of a hypothetical, simulated machine rather than a physical processor. This intermediate code can then be executed on any platform that has a

Downloaded from www.pdf.tube

parse tree, and peephole optimization improves efficiency by examining small instruction windows. However, type checking is a semantic analysis task, not syntactic. Furthermore, LALR parsers are more powerful than SLR parsers, as they use more refined lookahead information.

170. (D) Complex stack PDA.
Regular languages are the simplest class in the Chomsky Hierarchy and can be described by regular expressions or recognized by both deterministic and non-deterministic finite automata. These models are equivalent in power. However, a complex pushdown automaton with a stack is specifically designed for more powerful context-free languages. While it could recognize regular languages, it is not a standard description.

171. (A) Unchanged assignments.
In data-flow analysis, the concept of reaching definitions is used to track where each variable assignment in a program might still be valid without being overwritten. This information is crucial for various optimizations, such as constant propagation or identifying where a specific value can be safely used. It helps the compiler understand how data values propagate through the control-flow graph.

172. (A) Regular, Context-Free, Context-Sensitive, Unrestricted
The Chomsky Hierarchy organizes formal languages from most restrictive to most powerful. Regular languages are at the bottom, followed by context-free languages which introduce nested structures. Context-sensitive languages add dependencies based on surrounding symbols. Finally, unrestricted languages are the most general class, capable of describing any process that can be computed by a Turing machine. This reflects increasing grammar complexity.

173. (A) Labels and jump code.
When an intermediate code generator translates conditional control flow statements like if-then-else, it primarily uses labels and jump instructions. The generated code includes a conditional branch that tests the boolean expression and then jumps to the appropriate label for either the then or else block. This structure allows the program to skip sections of code based on results.

174. (A) Aid predictive choice.
Left-factoring is a grammar transformation technique used to prepare a grammar for predictive parsing methods like LL(1). It involves rewriting productions that share a common prefix into a new form that delays the decision between them. By removing this ambiguity, the parser can use its lookahead symbol to make a unique and correct choice between rules without backtracking.

175. (A) Optimization phase.
A simple interpreter's primary goal is to execute the source program as directly as possible, usually by traversing a syntax tree or processing tokens one by one. Because the emphasis is on immediate execution and flexibility, an elaborate optimization phase is often omitted. In contrast, compilers frequently include complex optimization stages to maximize the performance of the final machine code.

176. (A) Ambiguity check.
While many properties of context-free languages, such as membership and emptiness, can be determined by known algorithms, checking if a given context-free grammar is ambiguous is undecidable. This means there is no general procedure that can examine any context-free grammar and always correctly decide if it can produce more than one parse tree. This is a significant consideration.

177. (A) {a^n b^n | n ≥ 0}
A grammar with the production S → aSb | ε generates strings where each 'a' is matched by a corresponding 'b' at the end. The recursive rule allows for any number of these balanced pairs to be nested, while the empty string rule provides a base case for ending the derivation. This results in strings with 'a's followed by equal 'b's.

NODIA

178. (A) Record of the caller.
The control link, or dynamic link, is a pointer stored in an activation record that points back to the activation record of the calling procedure. This link is essential for managing the call stack at runtime. When a procedure finishes execution, the program uses this link to correctly restore the stack and return control to the caller, maintaining execution context.

in class P. In contrast, the general satisfiability problem is intractable. The halting problem is the most famous example of an undecidable problem. Finally, the set of strings accepted by a DFA perfectly exemplifies regular languages.

199. (A) Reducible stack portion.
In a shift-reduce parser, a handle refers to a specific portion of the symbols currently on the stack that can be reduced using a grammar production rule. Identifying the handle correctly is the central task of the parser. Once found, the handle is replaced with the corresponding non-terminal, representing a step in building the program's structural representation from tokens upward.

200. (A) Reclaim heap memory.
Garbage collection is an automated memory management mechanism used by run-time systems to reclaim memory in the heap that is no longer being used by the program. By identifying objects that are unreachable from current execution, the collector can safely free up that space for future allocations. This process prevents memory leaks and simplifies the developer's task by removing manual management.

201. (A) Both A and R are true and R is the correct explanation of A.
The language consisting of strings with a prime number of symbols is a classic example of a non-regular language. This can be rigorously proven using the Pumping Lemma for regular languages. By showing that any sufficiently long string in the language can be pumped to create a new string with a length that is composite rather than prime, a contradiction is established.

202. (A) Left factoring.
Left factoring is a fundamental grammar transformation technique used to prepare a grammar for predictive parsing methods like LL(1). It involves rewriting rules that share a common prefix to delay the choice between them until the next symbol is seen. This eliminates ambiguity and ensures that the parser can make a unique and correct decision using its lookahead symbol.

203. (A) Resolve cross-references.
The primary responsibility of a linker is to resolve cross-references between different pieces of compiled code. When a program is divided into multiple files, the linker combines these separate object files into a single cohesive executable. It ensures that every function call and variable reference correctly points to its definition, finalizing memory addresses and creating a complete program for execution.

204. (A) (i), (ii) and (iv) only
A standard Turing Machine is a powerful model of computation featuring an infinite tape and a head that can move both left and right. It recognizes the class of recursively enumerable languages. Despite the additional flexibility, a non-deterministic Turing machine is not computationally more powerful than a deterministic one; both models can solve exactly the same set of problems reliably.

205. (A) 2n - 1
In a grammar that is in Chomsky Normal Form, every derivation of a string of length 'n' requires exactly 2n - 1 steps. This is a direct consequence of the CNF rules, where each step either produces a single terminal symbol or splits a non-terminal into two new non-terminals. This predictable relationship is very useful for analyzing time complexity in parsing algorithms.

206. (A) Loop invariant motion.
Loop-invariant code motion is a classic example of a global optimization technique. It involves analyzing the entire control-flow graph of a procedure to identify calculations within a loop that produce the same result every iteration. By moving these calculations outside the loop, the compiler can significantly improve performance. This broad analysis sets global optimization apart from local techniques in small blocks.

207. (A) Context-Free Language.
A language that can be accepted by a pushdown automaton is classified as a context-free language. These languages are more expressive than regular languages, which are handled by finite automata. The addition of a stack allows the PDA to store and retrieve information, giving it memory to recognize structures like balanced parentheses and nested logical blocks common in programming language syntax.

Click Here to Buy book on Amazon

208. (A) Semantic associations.
The purpose of an attribute grammar is to provide a formal way to associate semantic information with the syntactic structures of a context-free grammar. By defining attributes and semantic rules for each production, the compiler can track important program details like data types and variable scopes. This approach integrates the analysis of meaning directly into the parsing process for efficient translation.

209. (A) Syntax Analysis.
Errors that involve the structural rules of the programming language, such as a missing semicolon or unbalanced braces, are typically detected during the syntax analysis phase of compilation. The parser, which is responsible for building a representation like a parse tree, identifies these mistakes when it encounters a sequence of tokens that does not match any valid production rule in grammar.

210. (C) Handle all CFGs.
LR parsers are a sophisticated class of bottom-up parsers that scan input from left to right and use a state stack to manage the derivation. While they are powerful and can handle a wide variety of grammars, it is incorrect to say they can handle all context-free grammars. There are certain complex or ambiguous grammars that remain beyond their reach.

211. (A) May loop on rejection.
A fundamental property of recursively enumerable languages is that they are recognized by a Turing machine that is guaranteed to halt and accept for any string that belongs to the language. However, for strings that are not in the language, the machine may either halt and reject or enter an infinite loop. This semi-decidability distinguishes RE languages from recursive ones.

212. (A) Infinite logic loop.
Left-recursion is a significant problem for top-down parsing methods because it can lead to an infinite recursion loop. When a grammar rule starts with the same non-terminal it is defining, a top-down parser will repeatedly call the corresponding function without ever consuming a token from the input. This prevents progress through the code and eventually causes the program to fail.

213. (A) Optimization is mandatory.
While the lexical, syntax, and semantic analysis phases are essential for translating and validating a program, the code optimization phase is generally considered optional. A compiler can produce a logically correct and executable output without performing any optimizations. However, most modern compilers include extensive optimization stages to ensure that the final machine code is fast and efficient for the target hardware.

214. (A) Lexical -> Syntax -> Semantic
A typical compiler processes source code through a series of logical phases in a specific order. It starts with lexical analysis to identify tokens, then moves to syntax analysis to build a structural representation, and finally performs semantic analysis to check for logical and type errors. This sequential process ensures each stage has validated information from the previous one for translation.

NODIA

215. (A) (i), (ii), (iv), and (v) only
Run-time environments use a variety of strategies to manage program data and control. Stack allocation is ideal for nested lifetimes, while the heap handles dynamic objects. Global variables are typically stored in a static area. An activation record is the core structure for managing procedure calls. However, executable code for procedures is usually kept in a separate text segment, not stack.

216. (A) Accepting states set.
In the formal definition of a deterministic finite automaton, the component 'F' signifies the set of accepting states. These are the specific states in which the machine must end its computation for an input string to be considered recognized or accepted by the language. For any given DFA, 'F' is a subset of the total set of states 'Q'.

217. (A) Constant folding.
Constant folding is a compiler optimization technique where expressions with constant operands are evaluated during compilation rather than at run time. If the value of an expression can be determined beforehand, the compiler replaces the entire expression with its calculated result in the final code. This reduces the number

of instructions executed and improves speed by avoiding redundant operations during runtime.

218. (A) Both Statement I and Statement II are true.
A language that requires a balance between two different symbols is context-free because it can be recognized by a pushdown automaton using its stack to track counts. As long as the relationship only involves two groups, the machine can simply push symbols for one and pop them for the other. Additional groups can often be processed independently without violating stack limits.

219. (B) Restore caller state.
The dynamic link, also known as the control link, is a pointer stored in an activation record that points back to the activation record of the caller. Its primary purpose is to allow the program to restore the caller's execution environment and stack frame once the current procedure has finished. This ensures the dynamic call sequence is correctly maintained during execution.

220. (A) Represent nested logic.
The primary limitation of a regular grammar is its inability to represent nested structures or balance an arbitrary number of symbols. Because regular grammars are recognized by finite automata, which have limited memory, they cannot handle recursive, balanced patterns common in languages. This makes them unsuitable for describing full syntax, which requires context-free grammars and pushdown automata for recognition.

221. (B) P-2, Q-4, R-1, S-3
Complexity classes provide a framework for categorizing problems based on difficulty. Class P contains problems solvable in polynomial time, while Class NP includes problems where solutions can be verified quickly. Undecidable problems, like the halting problem, have no algorithmic solution. Finally, NP-Hard problems are at least as challenging as the most difficult problems in NP, representing significant barriers in computational efficiency.

222. (A) Synthesized attributes.
The main characteristic of an S-attributed syntax-directed definition is its exclusive use of synthesized attributes. In this model, the value of an attribute at any given node is computed solely from the attribute values of its children in the parse tree. This information flows upward from the leaves to the root, allowing for efficient evaluation during a single bottom-up pass.

223. (A) A stream of tokens.
The primary output of the lexical analysis phase is a structured stream of tokens. As the lexer scans the source code's character stream, it identifies meaningful sequences called lexemes and categorizes them into tokens like keywords and identifiers. This token stream serves as the essential input for syntax analysis, which uses them to build a structural representation of the program.

224. (A) Equivalent DFA exists.
The equivalence between non-deterministic and deterministic finite automata implies that for every language that can be recognized by an NFA, there is an equivalent DFA that recognizes exactly the same set of strings. This fundamental result shows that while non-determinism might make design easier, it does not add additional power in terms of the class of languages recognized.

225. (A) Pass-by-name.
A thunk is an implementation technique closely associated with the pass-by-name parameter-passing mechanism. It is a small piece of code that wraps the expression passed as an argument. Each time the procedure accesses the formal parameter, the thunk is called to re-evaluate the expression in the caller's context. This ensures the parameter always reflects the most current expression value.

226. (A) Terminals and rules.
A context-free grammar is a formal system used to describe the syntax of a language. It consists of a finite set of terminal symbols, which are the basic elements, and non-terminal symbols representing categories. The core is a set of production rules defining how non-terminals can be replaced by symbol sequences, starting from a designated start symbol to generate valid strings.

227. (A) Both A and R are true and R is the correct explanation of A.
A Universal Turing Machine is a powerful theoretical model that can compute any function that is computable by any other machine. It

Downloaded from www.pdf.tube

achieves this by taking an encoded description of another Turing Machine and its input string and then simulating the step-by-step behavior of that specific machine. This confirms that a single general-purpose machine is capable of any computation.

228. (A) Reduce.
In bottom-up parsing, the action of using a grammar rule to replace a sequence of symbols on the parser's stack with a single non-terminal is called a reduce action. This occurs when the parser identifies a handle on the stack that matches the right-hand side of a rule. By performing this reduction, the parser builds the structural representation upward.

229. (A) Source code text.
An activation record is designed to hold all the necessary information for a single execution of a procedure, including local variables, return addresses, and parameters. However, the original source code text of the program is not stored within these records. Instead, the code is typically kept in a separate, read-only memory segment during the program's runtime for efficiency.

230. (A) Loop-invariant motion.
Loop-invariant code motion is an optimization technique that moves calculations whose results do not change during loop execution from inside the loop to a position just before it. By performing these computations only once instead of across every iteration, the compiler can significantly reduce total instructions executed. This leads to faster performance, especially for loops executed many times.

231. (A) Context-free.
The language consisting of strings where either the count of 'a's matches 'b's or 'b's matches 'c's is context-free. This is because the union of two context-free languages, each tracking one of the two possible balances, is itself a context-free language. A pushdown automaton can non-deterministically choose which balance to check using its stack, making the set recognizable.

232. (A) R, P, Q, S
The development of parsing algorithms has progressed from simple recursive descent methods to more sophisticated techniques. LL(1) parsing was a key early predictive method, followed by the development of LR(0) and its more powerful variants. LALR(1) parsing was developed later as a practical optimization to balance parsing power and memory efficiency. This chronological progression reflects increasing complexity in the field.

233. (A) Classify usability.
Distinguishing between tractable and intractable problems is essential for classifying the practical usability of computational solutions. While theoretically solvable, intractable problems require time that grows so quickly with input size that they are effectively impossible for real-world instances. This classification helps researchers focus on finding efficient algorithms or approximate solutions, ensuring the practical limits of computers are understood.

234. (A) `t1 = i * width; t2 = y + t1; x = *t2`
In three-address code, an array access like `x = y[i]` is broken down into multiple steps. The compiler first calculates the offset by multiplying the index by element width, then adds this to the base address of the array to find the exact memory location. Finally, it uses a load instruction to move the value into the target variable.

235. (B) Length $|\alpha| \leq |\beta|$.
A key feature of every production rule in a context-sensitive grammar is that the length of the string on the left side of the transformation must be less than or equal to the length on the right side. This non-contracting property ensures that strings do not shrink during derivation, which is a fundamental requirement for recognition by linear bounded automata.

236. (A) Dead code removal.
Dead code removal is an optimization technique where the compiler identifies and eliminates instructions that have no effect on the program's final output. If a variable is assigned a value but that value is never used later in any execution path, the assignment is unnecessary. Removing such redundant code reduces the size of the final executable and improves overall performance.

237. (B) (ii) and (iv) only
In compiler design, the symbol table is the central repository for storing identifier information

across all phases, and a Directed Acyclic Graph provides a compact way to represent common subexpressions. However, syntax analysis must precede semantic analysis to establish structure first. Type checking is a semantic task, not syntactic, and lexical analysis's role is identifying tokens, not syntax errors.

238. (A) Minimize DFA states.
The Myhill-Nerode theorem provides a mathematical basis for determining the minimal number of states required for any deterministic finite automaton to recognize a specific regular language. By defining an equivalence relationship between strings based on future behavior, the theorem allows researchers to prove whether a DFA is as small as possible or if further state reduction is achievable.

239. (A) Association ambiguity.
The dangling else problem refers to a specific type of association ambiguity that occurs in grammars with nested conditional statements. When an else clause appears after multiple if statements, the grammar rules may not clearly specify which if the else belongs to, potentially leading to multiple parse trees. Compilers typically resolve this by associating the else with the most recent unmatched if.

240. (C) Statement I is true, but Statement II is false.
The class of recursive languages is closed under complementation, meaning that if a language is fully decidable, its complement is also decidable. However, the same does not hold for recursively enumerable languages. The complement of an RE language is not necessarily recursively enumerable itself. This distinction illustrates the difference between machines that always halt and those that only halt for strings within.

241. (A) LALR(1) parser.
YACC is a classic parser generator tool that typically produces an LALR(1) parser from a given context-free grammar. This type of parser offers an excellent balance between expressive power and memory efficiency, making it suitable for a wide range of programming languages. By automatically generating parsing tables and state machine code, YACC simplifies the task of implementing syntax analysis.

242. (A) Stack memory usage.
The fundamental mechanism that allows a pushdown automaton to recognize context-free languages is its use of an external stack. This last-in, first-out data structure provides the machine with the necessary memory to handle nested and balanced structures that finite automata cannot manage. By pushing and popping symbols, the PDA can track counts and relationships across different parts of an input string.

243. (A) Undecidable.
The problem of determining if an arbitrary Turing machine will ever write a specific symbol on its tape is undecidable. This question is closely related to the halting problem, as determining any specific future action of a general-purpose machine over an infinite time horizon is impossible to solve with a general algorithm. No procedure can reliably predict if that specific action will occur.

244. (A) Parent/Left siblings.
A syntax-directed definition is categorized as L-attributed if the attributes of any symbol on the right-hand side of a production rule depend only on information from its parent node or from its siblings located to its left. This restriction ensures all attribute values are available during a single, left-to-right, top-down traversal of the parse tree, making them efficient for implementation.

245. (A) Directed Acyclic Graph (DAG).
A Directed Acyclic Graph is an intermediate representation that is particularly effective at identifying and managing common subexpressions within a program. By allowing multiple nodes to point to the same child, a DAG can represent a single calculation whose result is used in several different parts of an expression. This compact structure reduces redundancy and simplifies various optimization processes during compilation.

246. (A) Set theory flaw.
In the theory of computation and mathematical logic, Russell's Paradox demonstrates a fundamental flaw in naive set theory. By showing that the concept of the set of all sets that do not contain themselves leads to an unavoidable logical contradiction, it proved that early definitions were inconsistent. This discovery necessitated

Click Here to Buy book on Amazon

281. (A) To check the program for semantic correctness and perform type-checking.
Semantic analysis is a vital phase of compilation that goes beyond the basic structure of the code to ensure it makes logical sense. Its primary functions include verifying the semantic correctness of the program and performing rigorous type-checking. This ensures that all operations are performed on compatible data types and that the program adheres to the language's specific rules.

282. (B) It is computed from the attributes of its children nodes.
Synthesized attributes are a key component of syntax-directed definitions used during semantic analysis. The value of such an attribute at a specific node in the parse tree is derived directly from the attribute values of its children nodes. This mechanism allows information to flow upward through the tree structure, aggregating data from simpler components to form complex semantic information.

283. (B) The relationships and evaluation order among attributes.
A dependency graph is a graphical tool used by compilers to map out the relationships between various semantic attributes in a syntax-directed definition. By identifying which attributes rely on the values of others, the graph helps determine a valid evaluation order. This ensures that every attribute is calculated only after all the necessary information it requires has been properly determined.

284. (A) uses only synthesized attributes.
An S-attributed definition is a specialized category of syntax-directed definition where all semantic rules rely exclusively on synthesized attributes. Because information in these definitions only flows from the children nodes upward to the parent nodes, the attributes can be efficiently evaluated during a single bottom-up pass of the parse tree. This simplicity makes them useful for many compilation tasks.

285. (C) An SDD is a grammar augmented with semantic rules, which is a formal way to express an attribute grammar.
The relationship between an attribute grammar and a syntax-directed definition is one of theory and application. An attribute grammar provides the formal framework for associating semantic rules with a context-free grammar. A syntax-directed definition is the practical implementation of this framework, using these rules to specify how attributes should be calculated and propagated throughout the structural representation.

286. (D) Information like local variables, parameters, and saved machine status.
An activation record, also known as a stack frame, is a block of memory managed at run-time to store all information necessary for a single procedure call. It typically includes the procedure's local variables, the parameters passed to it, and the saved machine status, such as the return address. This structure ensures that each function execution remains isolated and manageable.

287. (C) The nested, LIFO nature of procedure calls and returns.
The use of a stack for allocating activation records is a fundamental design choice in most procedural programming languages. This strategy directly supports the last-in, first-out nature of procedure calls and returns, where the most recently called function is the first to complete. This synchronization between memory management and execution flow allows for efficient handling of nested and recursive calls.

288. (B) The sequence of procedure activations.
An activation tree provides a clear visualization of how procedures are invoked and completed during a program's execution. Each node in the tree represents a specific instance of a procedure being called. The tree structure effectively captures the hierarchical and temporal relationships between different activations, illustrating the sequence of calls and the return of control as logic progresses.

289. (D) To manage storage based on information about identifiers.
The symbol table is a central data structure that persists throughout the compilation and execution processes. During run-time, the system relies on the symbol table to manage memory allocation and data access. By storing detailed information about identifiers, such as their types and scopes, the table allows the system to

NODIA

regular languages. While one uses a generative approach involving production rules and the other utilizes an algebraic notation for pattern matching, both describe the same set of string patterns. This equivalence is a foundational concept that links different descriptive paradigms in automata theory.

256. (B) The application of the Pumping Lemma.
In formal language theory, the Pumping Lemma serves as a vital mathematical tool used to determine the boundaries of language classes. Its primary application is to demonstrate that a specific language does not belong to the regular class. By showing that a language fails the structural requirements for pumping strings, theorists can prove it lacks the necessary finite state characteristics.

257. (D) The intersection of L1 and L2.
Regular languages possess specific closure properties, meaning that applying certain operations to them results in another language within the same class. Intersection is one such operation; if two languages are proven to be regular, their common strings also form a regular language. This property is essential for modular design, allowing complex recognizers to be constructed from simpler, verified regular components.

258. (A) To use regular language principles to identify tokens in source code.
Lexical analysis serves as the initial phase of the compilation process, focusing on the basic building blocks of source code. It utilizes the principles of regular languages and regular expressions to scan the input text and identify specific patterns. These patterns are then categorized and transformed into a continuous stream of tokens, which are passed forward for syntax analysis.

259. (C) To prove a language is non-regular by showing a contradiction.
The Pumping Lemma is fundamentally a tool for proof by contradiction in automata theory. To utilize it, one assumes a language is regular and then identifies a sufficiently long string within that language. By demonstrating that any attempt to pump or repeat a segment of this string results in a word outside the language, the initial regularity assumption is invalidated.

260. (A) Sufficiently long strings must have a repetitive structure that allows for pumping.
The Pumping Lemma implies that all regular languages possess a specific repetitive structural property for strings exceeding a certain length. This means that any sufficiently long string within a regular language must contain a middle section that can be repeated indefinitely while remaining within the language. This inherent predictability is a direct consequence of the finite states in the automaton.

261. (D) A stack
A pushdown automaton is essentially a finite state machine enhanced with an internal stack. This additional component provides the machine with a form of unbounded memory, allowing it to keep track of information in a last-in, first-out manner. This capability is crucial for recognizing context-free languages, such as those containing balanced parentheses, which cannot be handled by finite automata.

262. (B) A string in the language can have more than one corresponding parse tree.
Ambiguity in a context-free grammar occurs when a single string belonging to the language can be derived in multiple ways, resulting in more than one distinct parse tree. This characteristic is generally viewed as a flaw in programming language design because it leads to uncertainty in how code should be interpreted. Compilers require a unique structural representation to ensure consistent execution.

263. (C) Context-Free Grammars (CFGs)
In the study of formal languages, non-deterministic pushdown automata and context-free grammars are established as equivalent models of representation and recognition. This means that for every language generated by a context-free grammar, there exists a corresponding non-deterministic pushdown automaton that accepts it. This equivalence parallels the relationship seen between finite automata and regular grammars in the Chomsky hierarchy.

NODIA

the development of more rigorous axiomatic set theories for a consistent foundation.

247. (A) P-3, Q-4, R-1, S-2
Each compiler task is associated with specific concepts. Parsing relies on context-free grammars to understand structure, while semantic analysis uses attribute grammars to handle meaning. Lexical analysis is driven by regular expressions to identify tokens from characters. Finally, runtime memory management uses the stack and heap to organize data and control flow as the program executes, ensuring correct translation.

248. (A) Execution order.
The control flow of a program refers to the specific order in which individual statements, instructions, or function calls are executed during its runtime. It is a fundamental concept describing how the program moves from one part of the code to another based on logic and results of conditional tests. Analyzing control flow is essential for understanding behavior and for performing optimizations.

249. (A) Multiple derivations.
A context-free grammar is classified as ambiguous if there exists at least one string in the language it generates that can be associated with more than one distinct derivation, resulting in multiple different parse trees. This ambiguity can lead to different structural interpretations, which is problematic. Compilers must resolve these, often through precedence and associativity rules, to ensure a single well-defined meaning.

250. (A) Type-0 Grammar
The Type-0 grammar, also known as an unrestricted grammar, is the most powerful class in the Chomsky Hierarchy. It has no limitations on its production rules, allowing any sequence of symbols to be replaced by any other sequence. This absolute flexibility allows Type-0 grammars to generate the most general class of languages, known as recursively enumerable languages, equivalent to Turing machines.

251. (B) It allows for flexible automaton design with the assurance of deterministic implementability.
The equivalence between deterministic and non-deterministic finite automata is a core principle in formal language theory. It ensures that while non-deterministic machines offer more flexibility and ease during the initial design phase, they do not exceed the computational limits of deterministic models. Consequently, every regular language designed non-deterministically can be successfully implemented using deterministic hardware or software systems.

252. (A) Pushdown Automaton (PDA)
Regular languages are represented through multiple equivalent computational and algebraic models. These include deterministic and non-deterministic finite automata, regular expressions, and regular grammars. However, pushdown automata are not part of this specific group. Instead, pushdown automata are specifically designed to recognize context-free languages, which occupy a higher level of complexity within the standard Chomsky hierarchy of formal languages.

253. (C) have multiple transition possibilities for a single input symbol.
The structural design of a non-deterministic finite automaton allows for greater flexibility compared to a deterministic model. Specifically, it permits the machine to have several potential transitions for a given combination of a current state and an input symbol. It also allows transitions to occur on empty strings, whereas deterministic machines are restricted to exactly one transition per input.

254. (D) provide a robust and interconnected theoretical framework.
Formal language theory relies on several equivalent models, such as automata, grammars, and algebraic expressions, to define regular languages. These various representations are significant because they collectively establish a robust and interconnected theoretical framework. This connectivity allows researchers and developers to choose the most convenient representation for a specific task while maintaining consistent mathematical properties across all chosen models.

255. (C) They are two distinct and equivalent methods for describing the same class of languages.
Regular grammars and regular expressions represent two distinct yet mathematically equivalent methods for defining the class of

Downloaded from www.pdf.tube

At the highest level, unrestricted grammars are used to generate the class of recursively enumerable languages. These languages are recognized by Turing Machines, which represent the most powerful computational model, capable of simulating any algorithm and managing the most complex structures.

273. (D) halt on a particular input string.
The Halting Problem is a cornerstone of theoretical computer science, famously proven to be undecidable by Alan Turing. It focuses on the fundamental question of whether a specific Turing Machine will eventually finish its computation or continue running forever when given a particular input string. This problem demonstrates that there are inherent limits to what can be predicted through algorithmic processes.

274. (A) As a tool to demonstrate the unsolvability of other problems.
The Post Correspondence Problem is a classic undecidable problem involving string matching in sequences. In computability theory, its primary value lies in its use as a baseline for proving the unsolvability of other problems. By reducing this problem to new, untested problems, researchers can demonstrate that the new problems are also undecidable, particularly those involving context-free language properties.

275. (C) The class of regular languages is a subset of the class of context-free languages.
The Chomsky Hierarchy illustrates a nested relationship between different classes of formal languages, where each level includes all the levels below it. Under this structure, the class of regular languages is a proper subset of the class of context-free languages. This implies that any language recognizable by a finite automaton can also be recognized by the more powerful pushdown automaton.

276. (A) Top-down parsers build the parse tree from the root, while bottom-up parsers build it from the leaves.
Top-down and bottom-up parsing represent two inverse strategies for constructing a hierarchical representation of a program's structure. Top-down methods begin at the grammar's start symbol and work toward the input tokens at the leaves. Conversely, bottom-up methods start with the individual tokens and perform a series of reductions to move upward until they reach the root of the tree.

277. (D) To resolve ambiguities in how expressions should be structured.
Rules of precedence and associativity are essential for resolving structural ambiguities that often occur during the syntax analysis phase of a compiler. These rules provide a clear framework for determining the correct order of operations in complex mathematical or logical expressions. Without these guidelines, a parser might generate multiple conflicting parse trees for the same sequence of tokens.

278. (B) LR Parser
Syntax analysis utilizes different parsing algorithms depending on the structure of the grammar and the direction of tree construction. The LR parser, along with its variant the LALR(1) parser, is a prominent example of a bottom-up parsing method. These parsers are highly efficient and are commonly used in modern compiler tools because they can handle a wide variety of grammars.

NODIA

279. (C) To make a grammar suitable for a top-down parsing algorithm.
Grammar transformations are often necessary to adapt a language's specification for use with specific parsing algorithms. For instance, eliminating left recursion is a critical step for making a grammar compatible with top-down parsing methods like recursive descent or LL(1). Without such modifications, the parser might enter an infinite loop or fail to accurately predict the correct production to apply.

280. (D) It takes a stream of tokens as input and produces a hierarchical syntactic structure as output.
The syntax analysis phase acts as a bridge between the linear scanning of source code and the structural understanding of a program. It takes a stream of tokens, generated during lexical analysis, as its primary input. The output is a hierarchical representation, such as a parse tree, which reflects the syntactic structure of the code according to the language's grammar.

264. (D) For standardization and to improve parsing efficiency.
Converting context-free grammars into standardized formats like Chomsky Normal Form or Greibach Normal Form is a standard practice in compiler design. These transformations are performed to simplify the grammar's structure, which in turn facilitates the creation of more efficient parsing algorithms. Standardizing the rules makes it easier to implement deterministic recognition and ensures consistent behavior across different parsing tools.

265. (A) Non-deterministic PDAs can recognize a broader class of languages than deterministic ones.
In the context of pushdown automata, there is a distinct difference in computational power between deterministic and non-deterministic versions. Unlike finite automata, where both types are equivalent, non-deterministic pushdown automata can recognize a wider range of languages than their deterministic counterparts. This makes them the primary model used for representing the full class of context-free languages in formal theory.

266. (B) That any function computable by an algorithmic process can be computed by a Turing Machine.
The Church-Turing Thesis is a fundamental postulate in computer science regarding the nature of effective calculability. It suggests that any function or problem that can be solved by following a clear, algorithmic process can also be computed by a Turing Machine. This thesis establishes the Turing Machine as the definitive theoretical benchmark for what is considered a computable process.

267. (B) A Turing Machine exists that halts on all possible input strings, either accepting or rejecting them.
A recursive language is defined by the existence of a Turing Machine that serves as a total decider for that language. This means the machine is guaranteed to halt for every possible input string, regardless of whether the string belongs to the language. Upon halting, the machine provides a definitive answer by either accepting or rejecting the input string without looping.

268. (A) It highlights the robustness of the Turing Machine as a universal model of computation.
Various enhancements to the standard Turing Machine, such as adding multiple tapes or introducing non-determinism, do not actually increase its fundamental computational power. The fact that these variations can all be simulated by a basic machine highlights the robustness of the model. This consistency reinforces the Turing Machine's status as a universal and reliable standard for defining absolute computability.

269. (C) a modern stored-program computer that can execute any valid program.
The Universal Turing Machine is a specialized model capable of simulating any other Turing Machine when provided with its description and input. This flexibility makes it a direct conceptual predecessor to the modern stored-program computer. Just as a physical computer can execute any program loaded into its memory, the Universal Turing Machine can perform any computation that any specific machine can.

270. (D) A TM recognizing a recursively enumerable language might not halt on strings outside the language.
The distinction between recursive and recursively enumerable languages lies in their halting behavior on strings not contained within the language. While a Turing Machine for a recursive language must halt on all inputs, a machine for a recursively enumerable language is only required to halt and accept valid strings. For strings outside the language, the machine may potentially enter an infinite loop.

271. (C) No algorithm can exist to solve all instances of the problem correctly.
In the field of computability, an unsolvable or undecidable problem is one for which a general algorithm cannot be constructed to provide a correct answer for all instances. This means that no Turing Machine can exist that is guaranteed to halt and give a valid response for every possible case of the problem, regardless of its specific difficulty or complexity.

272. (B) Unrestricted languages
The Chomsky Hierarchy organizes formal grammars and their corresponding automata into four distinct levels of increasing expressive power.

Click Here to Buy book on Amazon

correctly map variable names to their specific locations in memory.

290. (A) The specific mechanisms used for passing parameters.
The run-time system is responsible for several critical tasks that occur while a program is actually executing. One of these primary responsibilities is managing the specific mechanisms used for parameter passing between procedures. Whether using pass-by-value or pass-by-reference, the system ensures that data is correctly transferred and accessible within the appropriate activation records during the lifespan of functions.

291. (C) It allows the compiler to be more easily adapted to different target machines.
A machine-independent intermediate representation serves as a versatile bridge between the source language and the target machine's hardware. One of its main advantages is that it allows the compiler to be more easily retargeted. By isolating the front-end analysis from the back-end code generation, developers can adapt the compiler to different architectures without rewriting the entire analysis phase.

292. (A) By using a combination of labels and conditional jump instructions.
Translating high-level control flow structures into low-level intermediate code requires breaking down complex logic into simpler, jump-based operations. For an if-then-else statement, the compiler typically generates a sequence of labels and conditional jump instructions. These jumps allow the execution path to skip or enter specific blocks of code based on the truth value of the initial conditional expression.

293. (D) The front-end analysis and the back-end synthesis.
The phase of intermediate code generation occupies a pivotal position within a compiler's architecture. It functions as the connection point between the front-end analysis phases and the back-end synthesis phases. By producing a machine-independent representation, this phase ensures that the structural and semantic information gathered during analysis can be effectively transformed into optimized instructions for the target machine.

294. (B) Source code
Intermediate representations are designed to be machine-independent formats that facilitate optimization and translation. Common examples include three-address code, syntax trees, and stack-based code. While source code is the original input to the compiler, it is not considered an intermediate representation itself. Instead, it is the starting point that is systematically transformed into more manageable, low-level internal formats.

295. (C) Generating code to set up an activation record and pass parameters.
The translation of a procedure call into intermediate code involves several systematic steps to prepare for execution. These actions include generating specific code to set up a new activation record on the stack and managing the passing of parameters to the called function. This process ensures that the program state is correctly preserved and that control flow transitions occur smoothly.

296. (D) To improve the run-time performance or memory usage of the program.
Code optimization is a crucial phase dedicated to enhancing the efficiency of the program before final machine code is generated. Its primary objective is to modify the intermediate representation so that the resulting software executes faster or consumes fewer resources, such as memory. This involves a variety of techniques that streamline the code without changing its original intended logic.

NODIA

297. (A) By examining a small, moving sequence of instructions and making local improvements.
Peep-hole optimization is a relatively simple yet highly effective technique used to improve the quality of target code. It involves scanning a small, continuous window of instructions to identify patterns that can be replaced with more efficient sequences. This local analysis allows for immediate improvements, such as removing redundant operations or replacing slow instructions with faster, equivalent alternatives.

298. (B) Understanding how data values are propagated through the program for optimization.
Data-flow analysis is a fundamental prerequisite for performing advanced code optimizations.

This process involves gathering detailed information about how data values move and change throughout a program's execution paths. By understanding these propagation patterns, a compiler can identify opportunities for optimization, such as removing unused code or identifying constants that can be calculated at compile time rather than run-time.

299. (B) Global optimization operates on larger program units like functions, while local optimization works on basic blocks.
The scope of analysis is the primary factor that distinguishes global optimization from local optimization. Local optimization is restricted to basic blocks, which are simple, straight-line sequences of code. In contrast, global optimization takes a broader view, operating across entire functions or procedures. This larger scope allows the compiler to make improvements that involve complex control-flow and data-flow interactions.

300. (A) Code Generation
Instruction scheduling is a critical component of the code generation phase, which is the final stage of the compilation process. This technique involves reordering the target machine's instructions to maximize execution efficiency and take full advantage of the processor's architecture. By carefully arranging the sequence of operations, the compiler can minimize delays and ensure the most effective hardware utilization.

Downloaded from www.pdf.tube

CHAPTER 09

Data Communication and Computer Networks

NODIA

1. Which of the following is considered a fundamental component of a data communication system?
(A) Algorithm
(B) Sender
(C) Compiler
(D) Programmer

2. In which mode of communication can data flow in both directions simultaneously?
(A) Full-duplex
(B) Simplex
(C) Half-duplex
(D) Multiplex

3. Which of the following statements about digital signals is correct?
(A) They are continuous waves with varying amplitude.
(B) They are less immune to noise compared to analog signals.
(C) They are discrete, time-varying signals.
(D) They cannot be transmitted over long distances.

4. A walkie-talkie is a classic example of which communication mode?
(A) Simplex
(B) Full-duplex
(C) Null-duplex
(D) Half-duplex

5. Match List I with List II regarding data communication concepts.
List I: (Concept)
P. Bandwidth
Q. Throughput
R. Latency
S. Simplex
List II: (Definition)
1. The time it takes for a message to completely arrive.
2. The actual measured rate of successful data transfer.
3. A unidirectional communication channel.
4. The theoretical range of frequencies a channel can pass.

Choose the correct option from those given below:
(A) P-4, Q-2, R-1, S-3
(B) P-2, Q-4, R-3, S-1
(C) P-4, Q-1, R-2, S-3
(D) P-3, Q-2, R-1, S-4

6. Given below are two statements regarding communication channels:
Statement I : A noiseless channel, as described by Nyquist, has a finite bit rate limit even without noise.
Statement II : A noisy channel's capacity, according to Shannon, is dependent on bandwidth and the signal-to-noise ratio.
Which of the above statements are correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

7. What is the primary purpose of modulation in data transmission?
(A) To convert a digital signal to another digital signal.
(B) To convert an analog signal to a digital signal.
(C) To reduce the bandwidth of the transmission.
(D) To encode digital data onto an analog carrier signal.

8. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : In a data communication

system, the medium is the physical path by which a message travels.

Reason (R) : Transmission media are always tangible, like wires and cables.

In light of the above statements, choose the most appropriate answer from the options given below:

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

9. Which technique combines several signals for transmission on a single medium?
(A) Modulation
(B) Demodulation
(C) Multiplexing
(D) Error correction

10. Which of the following is an example of an error handling mechanism?
(A) Cyclic Redundancy Check (CRC)
(B) Baseband transmission
(C) Amplitude Shift Keying (ASK)
(D) Time Division Multiplexing (TDM)

11. Consider the following statements about data transmission:
I. Baseband transmission uses the entire bandwidth for a single signal.
II. Broadband transmission divides the channel bandwidth among multiple signals.
III. Baseband transmission is typically used for digital signals over short distances.
IV. Broadband transmission uses analog signals.

Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

12. The term 'throughput' in data communication refers to the:
(A) Theoretical maximum data rate of a channel.
(B) Actual rate at which data is successfully transferred.
(C) Delay for the first bit to travel from source to destination.
(D) Number of errors per million bits transmitted.

13. Which of the following statements about network topologies is incorrect?
(A) In a star topology, all devices are connected to a central hub.
(B) A bus topology uses a single backbone cable to connect all devices.
(C) A mesh topology is the least reliable due to single points of failure.
(D) A ring topology connects devices in a circular fashion.

14. What type of network would typically be used to connect the offices of a company located in different cities?
(A) Local Area Network (LAN)
(B) Metropolitan Area Network (MAN)
(C) Wireless Network
(D) Wide Area Network (WAN)

15. Which of the following is NOT a type of computer network classified by geographical scope?
(A) Layered Area Network
(B) Wide Area Network
(C) Local Area Network
(D) Metropolitan Area Network

16. The concept of dividing network communication tasks into a vertical stack of distinct layers is known as:
(A) Network topology
(B) Layered architecture
(C) Switching technique
(D) Multiplexing

17. Consider the following statements about network models:
I. The OSI reference model is a practical protocol suite used for the internet.
II. The TCP/IP protocol suite is a theoretical model and not implemented.
III. The OSI model consists of seven layers.
IV. Layered architecture makes network design more modular.
V. The Physical layer is the lowest layer in the OSI model.

Which of the above statements are correct?

Click Here to Buy book on Amazon

(A) I, II, and III only
(B) III, IV, and V only
(C) I, III, and V only
(D) II, IV, and V only

18. What is the primary function of the OSI Physical Layer?
(A) Transmitting raw bits over a communication channel.
(B) Framing and error control.
(C) Routing packets across the network.
(D) Providing end-to-end communication services.

19. In the context of network addressing, which address is used to identify a process on a host?
(A) Physical address
(B) Logical address
(C) Port address
(D) Specific address

20. Match List I (OSI Layer) with List II (Primary Function).
List I: (OSI Layer)
P. Application Layer
Q. Transport Layer
R. Network Layer
S. Data Link Layer
List II: (Primary Function)
1. Framing and physical addressing.
2. Routing and logical addressing.
3. User interface and network services.
4. Process-to-process delivery and error control.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-2, S-1
(C) P-3, Q-2, R-4, S-1
(D) P-1, Q-4, R-2, S-3

21. A MAC address is an example of which type of address?
(A) Logical address
(B) Port address
(C) Physical address
(D) Specific address

22. Which switching technique involves dedicating a complete physical circuit between two nodes for the duration of their communication?
(A) Packet switching
(B) Circuit switching
(C) Message switching
(D) Frame switching

23. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The TCP/IP protocol suite is the de facto standard for the internet.
Reason (R) : The OSI model was developed after the TCP/IP suite and is more practical for implementation.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

24. The process of adding headers and trailers to data at each layer of a network model is called:
(A) Encapsulation
(B) Fragmentation
(C) Multiplexing
(D) Modulation

25. The function of dividing a stream of bits from the Network layer into manageable data units called frames is performed by which layer?
(A) Physical Layer
(B) Transport Layer
(C) Data Link Layer
(D) Session Layer

26. Consider the following statements about error control:
I. Error detection involves identifying that an error has occurred.
II. Error correction involves identifying and correcting the corrupted bits.
III. Parity check is a form of error correction.
IV. Error control is primarily a function of the Data Link and Transport layers.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and IV only
(C) III and IV only

Downloaded from www.pdf.tube

(B) It acts as the physical hardware of the mobile device.
(C) It is the operating system of the mobile device.
(D) It is a type of application software, like a web browser.

69. The Mobile IP protocol is designed to:
(A) Increase the speed of mobile data connections.
(B) Encrypt all data transmitted from a mobile device.
(C) Allow a mobile device to maintain its IP address while moving between different networks.
(D) Provide cloud storage for mobile devices.

70. Consider the following statements about mobile and wireless technologies:
I. Communication satellites can be used to relay signals for mobile communication.
II. A cellular topology divides a service area into smaller regions called cells.
III. Mobile Adhoc Networks (MANETs) require a fixed infrastructure.
IV. Wireless LANs (WLANs) are governed by standards like IEEE 802.11.
V. GPRS provides packet-switched data services on GSM networks.
Which of the above statements are correct?
(A) I, II, III, and IV
(B) I, III, and V
(C) I, II, IV, and V
(D) II, III, and V

71. A gateway in the context of mobile computing primarily serves to:
(A) Store user data securely.
(B) Manage the device's battery life.
(C) Provide the user interface for mobile applications.
(D) Connect a mobile network to other networks, translating protocols if necessary.

72. What does SMS stand for in mobile technology?
(A) Short Mobile System
(B) Short Message Service
(C) Secure Message Service
(D) Satellite Mobile Service

73. A key characteristic of Mobile Adhoc Networks (MANETs) is that they are:
(A) Centrally managed by a base station.
(B) Limited to a very short range, like a single room.
(C) Infrastructure-less and self-configuring.
(D) Primarily used for voice communication only.

74. Wireless Geolocation Systems are used to:
(A) Increase the data transmission speed of a wireless network.
(B) Secure wireless communications through encryption.
(C) Manage the allocation of radio frequencies.
(D) Determine the geographical position of a mobile device.

75. Match List I (Cloud Service Model) with List II (Description).
List I: (Cloud Service Model)
P. SaaS
Q. PaaS
R. IaaS
List II: (Description)
1. Provides virtualized computing resources like servers and storage over the internet.
2. Provides a platform allowing customers to develop, run, and manage applications without the complexity of building infrastructure.
3. Provides software applications over the internet on a subscription basis.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3
(B) P-3, Q-1, R-2
(C) P-3, Q-2, R-1
(D) P-2, Q-3, R-1

NODIA

76. In cloud computing, a cloud that is open for use by the general public is known as a:
(A) Private Cloud
(B) Public Cloud
(C) Hybrid Cloud
(D) Community Cloud

77. The technology that allows a single physical computer to run multiple operating systems or multiple instances of one OS simultaneously is called:

37. Which statement best describes classful addressing in IPv4?
(A) The address space is divided into blocks of variable size.
(B) It uses a subnet mask to determine the network portion of the address.
(C) The address space is divided into five classes: A, B, C, D, and E.
(D) It is the modern standard for IP address allocation.

38. The process of breaking an IPv4 datagram into smaller units to pass through a network with a smaller MTU is called:
(A) Checksum
(B) Fragmentation
(C) Encapsulation
(D) Aggregation

39. What is the primary motivation behind the development of IPv6?
(A) To increase overall data transfer speeds across the global internet.
(B) To simplify the core TCP protocol for improved network performance.
(C) To eliminate the functional requirement for the data link layer.
(D) To provide more addresses and overcome IPv4 address exhaustion.

40. Consider the following statements about IPv4 and IPv6:
I. IPv4 addresses are 32 bits long.
II. IPv6 addresses are 128 bits long.
III. Checksum is a mandatory field in the IPv6 header.
IV. Classful addressing is a feature of IPv6.
V. Fragmentation is handled by routers in IPv6.
Which of the above statements are incorrect?
(A) I and II only
(B) III, IV, and V only
(C) I, III, and IV only
(D) II and V only

41. What protocol is used for mapping a logical (IP) address to a physical (MAC) address?
(A) DNS
(B) UDP
(C) ARP
(D) RARP

42. When the sender and receiver of an IP packet are on the same network, the delivery is referred to as:
(A) Direct delivery
(B) Indirect delivery
(C) Routed delivery
(D) Proxied delivery

43. Which of the following is NOT a routing algorithm?
(A) Distance Vector
(B) Link State
(C) Path Vector
(D) Sliding Window

44. Which transport layer protocol provides connectionless, unreliable datagram service?
(A) TCP
(B) UDP
(C) SCTP
(D) HDLC

45. Arrange the following steps in the process of a three-way handshake for TCP connection establishment in the correct chronological order:
1. The server sends a SYN-ACK segment.
2. The client sends a SYN segment.
3. The client sends an ACK segment.
Choose the correct sequence:
(A) 1, 2, 3
(B) 2, 1, 3
(C) 2, 3, 1
(D) 3, 2, 1

46. What mechanism does TCP use for flow control?
(A) A token passing system
(B) A polling mechanism
(C) A fixed-rate transmission
(D) A sliding window

47. In the context of TCP, what is congestion control primarily concerned with?
(A) Preventing the network from being overwhelmed by excessive traffic.
(B) Preventing the sender from overwhelming the receiver's available buffer.
(C) Ensuring that all packets arrive at the destination in the correct sequence.
(D) Correcting bit-level errors that occur during

Click Here to Buy book on Amazon

109. What is a Backbone Network?
(A) A small network segment connecting a few devices.
(B) A wireless network with no fixed infrastructure.
(C) A centralized, high-capacity network that connects other smaller, independent networks.
(D) The physical cable used in a bus topology.

110. In classless addressing (CIDR), the network portion of an IP address is defined by:
(A) A variable-length prefix or netmask.
(B) The first octet of the address.
(C) A fixed class (A, B, or C).
(D) The total number of hosts on the network.

111. The IPv6 packet format has a fixed-size base header. What is the primary advantage of this design?
(A) It allows for more optional features to be included.
(B) It reduces the total size of every packet.
(C) It eliminates the need for a transport layer protocol.
(D) It simplifies and speeds up packet processing by routers.

112. If a host on one network wants to send a datagram to a host on another network, the delivery is:
(A) Local
(B) Indirect
(C) Direct
(D) Unicast

113. What is the primary difference between a worm and a virus?
(A) A worm is malicious, while a virus is not.
(B) A worm targets networks, while a virus targets individual computers.
(C) A virus requires a host program to replicate, while a worm is a standalone program that can propagate itself.
(D) A virus encrypts data, while a worm steals it.

114. In mobile computing, the process of transferring an active call or data session from one cell to another without interruption is known as:
(A) Roaming
(B) Geolocation
(C) Multiplexing
(D) Handoff

115. Consider the following statements about cloud deployment models:
I. A public cloud is owned by a cloud provider and offered to the general public.
II. A private cloud is operated solely for a single organization.
III. A hybrid cloud combines public and private clouds.
IV. A community cloud is shared by several organizations with common concerns.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

116. A device that operates at the physical layer and regenerates the signal to extend the transmission distance is a:
(A) Switch
(B) Repeater
(C) Router
(D) Bridge

117. What does the term cellular topology refer to?
(A) The star-like physical layout of a network.
(B) The practice of dividing a geographical service area into a grid of cells, each served by a base station.
(C) A network where every node is connected to every other node.
(D) The hierarchical structure of the internet backbone.

118. Which statement about the Internet of Things (IoT) is most accurate?
(A) IoT focuses on connecting everyday physical devices to the internet to collect and share data.
(B) IoT is primarily about social networking on mobile devices.
(C) IoT is a new standard for high-speed fiber optic communication.
(D) IoT is an alternative name for cloud computing.

NODIA

(B) A single key is used for both encryption and decryption.
(C) The key is shared between the sender and receiver.
(D) It is generally faster than secret-key algorithms.

59. What is the primary purpose of a digital signature in network security?
(A) To ensure the confidentiality of a message.
(B) To hide the message from unintended recipients.
(C) To encrypt the entire communication channel.
(D) To provide authentication, non-repudiation, and integrity.

60. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : A firewall acts as a barrier between a trusted internal network and an untrusted external network.
Reason (R) : Firewalls work by inspecting network traffic packets and permitting or blocking them based on a set of security rules.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

61. A Virtual Private Network (VPN) is used to:
(A) Create a secure, encrypted connection over a less secure network, like the internet.
(B) Physically isolate a network from all external connections.
(C) Manage domain name resolution.
(D) Detect and remove malware from a host system.

62. Which of the following protocols is used for accessing and managing mail on a remote server, allowing for folder management on the server side?
(A) SMTP
(B) IMAP
(C) POP
(D) TELNET

63. The TELNET protocol is primarily used for:
(A) Secure file transfer.
(B) Sending and receiving electronic mail.
(C) Providing a command-line interface to a remote computer.
(D) Browsing web pages securely.

64. What is the main difference between FTP and TELNET?
(A) FTP is secure, while TELNET is not.
(B) FTP uses TCP, while TELNET uses UDP.
(C) FTP is for email, while TELNET is for web browsing.
(D) FTP is for file transfer, while TELNET is for remote login.

65. Consider the following statements about email protocols:
I. SMTP is a push protocol used to send mail.
II. POP is a pull protocol used to retrieve mail.
III. IMAP allows a user to manage emails directly on the server.
IV. FTP is a protocol used for email architecture.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

66. In mobile technology, GSM and CDMA are two competing:
(A) Mobile operating systems.
(B) Physical layer technologies for cellular networks.
(C) Cloud computing service models.
(D) Application layer protocols.

67. Which of the following is a key component of the GSM architecture?
(A) Public Switched Telephone Network (PSTN)
(B) Domain Name System (DNS)
(C) Internet Service Provider (ISP)
(D) Base Station Subsystem (BSS)

68. What is the role of middleware in mobile computing?
(A) It provides services to applications beyond what is available from the operating system.

Downloaded from www.pdf.tube

88. Which network topology is characterized by its high reliability due to redundant paths between devices?
(A) Mesh
(B) Bus
(C) Star
(D) Ring

89. An IP address is an example of which type of address?
(A) Physical address
(B) Port address
(C) Logical address
(D) Specific address

90. Consider the following statements about the functions of OSI layers:
I. Framing is a function of the Data Link Layer.
II. Routing is a function of the Network Layer.
III. Congestion control is a function of the Transport Layer.
IV. Synchronization is a function of the Session Layer.
V. Encryption is a function of the Presentation Layer.
Which of the above statements are correct?
(A) I, II, III, IV, and V
(B) I, II, and III only
(C) I, II, IV, and V only
(D) II, III, and IV only

91. The ability of a station to transmit on a multiple-access channel is determined by a central controller in which scheme?
(A) CSMA/CD
(B) Token Passing
(C) ALOHA
(D) Polling

92. Class C addresses in the original classful IPv4 addressing scheme were intended for:
(A) Large organizations with many hosts.
(B) Networks with a small number of hosts.
(C) Multicasting purposes.
(D) Experimental use.

93. What is a key difference between TCP and SCTP?
(A) TCP is connection-oriented, while SCTP is connectionless.
(B) TCP provides error control, while SCTP does not.
(C) SCTP supports multi-streaming, while TCP does not.
(D) SCTP is an application-layer protocol, while TCP is a transport-layer protocol.

94. The process of mapping a domain name like 'www.example.com' to its corresponding IP address is known as:
(A) Reverse mapping
(B) Forward mapping
(C) Address resolution
(D) Name translation

95. Arrange the following OSI layers in the order that data passes through them during encapsulation at the sender's side (from top to bottom):
1. Network Layer
2. Application Layer
3. Data Link Layer
4. Transport Layer
Choose the correct sequence:
(A) 2, 4, 1, 3
(B) 2, 1, 4, 3
(C) 4, 2, 1, 3
(D) 1, 4, 3, 2

96. Which statement regarding POP (Post Office Protocol) is generally true?
(A) It keeps all emails stored on the server by default.
(B) It is designed to download emails to the local client and then delete them from the server.
(C) It provides advanced folder management on the server.
(D) It is used for sending emails from the client.

97. In public-key cryptography, if A wants to send an encrypted message to B, A encrypts the message using:
(A) A's public key
(B) B's public key
(C) A's private key
(D) B's private key

98. Which mobile technology standard is based on Time Division Multiple Access (TDMA)?
(A) GSM
(B) CDMA

NODIA

(A) Virtualization
(B) Multiplexing
(C) Multitasking
(D) Fragmentation

78. What does IaaS stand for in the context of cloud computing?
(A) Internet as a Service
(B) Integration as a Service
(C) Infrastructure as a Service
(D) Information as a Service

79. A Service Level Agreement (SLA) in cloud computing is a contract between:
(A) Two cloud users.
(B) The cloud provider and the government.
(C) A user and an application developer.
(D) A cloud service provider and a consumer.

80. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Virtualization is a foundational element of cloud computing.
Reason (R) : Virtualization allows for the efficient sharing of physical resources among multiple users and applications, enabling key cloud features like elasticity and resource pooling.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

81. What is the Internet of Things (IoT)?
(A) A comprehensive upgrade of the standard World Wide Web architecture.
(B) A network of physical objects using sensors to exchange data.
(C) A specialized security protocol for managing online financial transactions.
(D) A theoretical model for establishing future decentralized networks.

82. In cloud computing, the term 'Virtual Server' refers to:
(A) A software-based emulation of a physical server.
(B) A physical server located in a remote data center.
(C) A backup server used for disaster recovery.
(D) A server that does not have an operating system.

83. Which of the following is a primary concern for Database Storage in a cloud environment?
(A) User interface design
(B) Processor speed of the client device
(C) The physical location of the user
(D) Data consistency, availability, and security

84. In GPRS, what kind of service is provided?
(A) Packet-switched data service
(B) Circuit-switched voice service
(C) Analog television broadcasting
(D) Secure shell access

85. Consider the following statements about network security concepts:
I. Cryptography is the practice of secure communication in the presence of adversaries.
II. Steganography aims to conceal the very existence of a message.
III. A digital signature uses symmetric-key cryptography.
IV. A firewall can be implemented in either hardware or software.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and IV only
(C) III and IV only
(D) I, II, III, and IV

86. A telephone conversation is an example of which communication mode?
(A) Simplex
(B) Full-duplex
(C) Half-duplex
(D) Simplex and Half-duplex

87. Which of the following is NOT a guided transmission medium?
(A) Twisted-pair cable
(B) Coaxial cable
(C) Fiber-optic cable
(D) Atmosphere

Click Here to Buy book on Amazon

the transmission of a packet.

48. Which protocol is a transport layer protocol that provides services similar to TCP but with additional features like multi-homing and multi-streaming?
(A) UDP
(B) SCTP
(C) ARP
(D) ICMP

49. What is the purpose of the checksum field in the IPv4 datagram header?
(A) To ensure the packet reaches the correct destination.
(B) For error detection in the header.
(C) To control the flow of data.
(D) To identify the version of the IP protocol.

50. Match List I (Protocol) with List II (Primary Purpose).
List I: (Protocol)
P. TCP
Q. UDP
R. ARP
S. IP
List II: (Primary Purpose)
1. Connectionless datagram service at the transport layer.
2. Mapping logical addresses to physical addresses.
3. Logical addressing and routing.
4. Connection-oriented, reliable service at the transport layer.
Choose the correct option from those given below:
(A) P-4, Q-1, R-2, S-3
(B) P-1, Q-4, R-3, S-2
(C) P-4, Q-2, R-1, S-3
(D) P-2, Q-1, R-3, S-4

51. What does URL stand for in the context of the World Wide Web?
(A) Uniform Resource Link
(B) Universal Resolution Locator
(C) Uniform Resource Locator
(D) Universal Resource Link

52. The primary function of the Domain Name Service (DNS) is to:
(A) Transfer files between computers.
(B) Manage network security through firewalls.
(C) Provide a secure shell for remote login.
(D) Resolve human-readable domain names into IP addresses.

53. Which of the following statements about DNS resolution is correct?
(A) Mapping physical MAC addresses to logical IP addresses for routing.
(B) Mapping host names to IP addresses while excluding reverse lookups.
(C) Mapping names to addresses and addresses to names in the system.
(D) Resolving domain names through local machines without external servers.

54. In the context of electronic mail architecture, which protocol is primarily used for sending email from a client to a server?
(A) POP3
(B) SMTP
(C) IMAP
(D) FTP

NODIA

55. Which of the following is NOT a type of malware?
(A) Virus
(B) Worm
(C) Trojan Horse
(D) Firewall

56. What is the core principle of steganography?
(A) To hide the message's existence within another file or medium.
(B) To transform the message into an unreadable form using encryption.
(C) To generate a unique digital signature for user authentication.
(D) To establish a dedicated private network over public infrastructure.

57. Secret-key cryptography is also known as:
(A) Asymmetric cryptography
(B) Public-key cryptography
(C) Symmetric cryptography
(D) Hash-based cryptography

58. Which of the following is a characteristic of public-key algorithms?
(A) It uses two keys: a public key for encryption and a private key for decryption.

(C) GPRS
(D) MANET

99. Providing an entire application to a client, who then runs that application on the cloud infrastructure, is an example of:
(A) PaaS
(B) IaaS
(C) SaaS
(D) DaaS

100. Match List I (Term) with List II (Belongs to).
List I: (Term)
P. Virtual LAN (VLAN)
Q. Classful Addressing
R. Mobile IP
S. Digital Signature
List II: (Belongs to)
1. Network Security
2. IPv4
3. Mobile Communication
4. LAN Technology
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-3, R-2, S-1
(C) P-4, Q-2, R-3, S-1
(D) P-2, Q-4, R-1, S-3

101. Which of the following is NOT a component of a data communication system?
(A) Protocol
(B) Message
(C) Operating System
(D) Medium

102. In a noiseless channel with bandwidth 'B' and 'L' signal levels, what is the maximum bit rate according to the Nyquist Bit Rate formula?
(A) BitRate = B * log2(1 + SNR)
(B) BitRate = 2 * B * log2(L)
(C) BitRate = B * L
(D) BitRate = 2 * B + log2(L)

103. Broadband transmission is characterized by:
(A) Using digital signals only.
(B) Using the entire bandwidth of the medium for one channel.
(C) Being used only for short-distance communication.
(D) Using analog signals and FDM to carry multiple channels.

104. What is the most significant disadvantage of a bus topology?
(A) A break in the main cable can disable the entire network.
(B) It is very expensive to install.
(C) It requires a central controller, which is a single point of failure.
(D) Adding new devices is very complex.

105. A network connecting all the computers within a single university campus is best described as a:
(A) WAN
(B) MAN
(C) LAN
(D) PAN

106. Which OSI layer is responsible for translating, encrypting, and compressing data?
(A) Presentation Layer
(B) Application Layer
(C) Session Layer
(D) Transport Layer

NODIA

107. Which address uniquely identifies a device on a local area network?
(A) Port address
(B) IP address
(C) Domain name
(D) MAC address

108. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Flow control is needed to prevent a fast sender from overwhelming a slow receiver.
Reason (R) : The Sliding Window Protocol is a common method for implementing flow control at the Transport and Data Link layers.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

(D) I, II, III, and IV

27. The Sliding Window Protocol is a mechanism for:
(A) Error detection only
(B) Multiple access
(C) Physical addressing
(D) Flow control

28. Which multiple access method requires a station to listen to the channel before transmitting and to stop transmitting if a collision is detected?
(A) CSMA/CA
(B) CSMA/CD
(C) FDMA
(D) TDMA

29. Consider the following network devices:
I. Hub
II. Switch
III. Router
IV. Repeater
V. Bridge
Which of these devices operate primarily at the Data Link Layer of the OSI model?
(A) I and IV only
(B) II and V only
(C) III only
(D) II, III, and V only

30. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : Routers operate at the Network layer of the OSI model.
Reason (R) : Routers use logical (IP) addresses to make forwarding decisions for packets between different networks.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

31. In the context of multiple access, which method is primarily used in wireless networks to avoid collisions?
(A) CSMA/CA
(B) CSMA/CD
(C) Token Passing
(D) Polling

32. What is the primary purpose of a Virtual LAN (VLAN)?
(A) To physically segment a LAN into smaller networks.
(B) To logically segment a LAN into different broadcast domains.
(C) To increase the physical size of a LAN.
(D) To connect multiple different LANs over a WAN.

33. Which of the following is a channelization-based multiple access protocol?
(A) CSMA/CD
(B) CDMA
(C) Polling
(D) Reservation

34. The High-Level Data Link Control (HDLC) protocol is most closely associated with which layer of the OSI model?
(A) Data Link Layer
(B) Physical Layer
(C) Network Layer
(D) Transport Layer

35. Consider the following statements regarding multiple access methods:
I. In Polling, a primary device invites secondary devices to transmit one by one.
II. In Token Passing, a special packet called a token circulates among the nodes.
III. Reservation is a method where stations reserve a future time slot for transmission.
IV. CSMA/CD is a contention-based protocol.
Which of the above statements are correct?
(A) I and II only
(B) III and IV only
(C) I, II, and IV only
(D) I, II, III, and IV

36. In IPv4 addressing, what is the size of an address?
(A) 16 bits
(B) 64 bits
(C) 128 bits
(D) 32 bits

Downloaded from www.pdf.tube

119. A user interacts with a web-based email client (like Gmail) in a browser. This entire service, including the software and the backend infrastructure, is an example of:
(A) Infrastructure as a Service (IaaS)
(B) Platform as a Service (PaaS)
(C) Software as a Service (SaaS)
(D) Network as a Service (NaaS)

120. Consider the following statements regarding network addressing:
I. A physical address (MAC) is used for node-to-node delivery on the same network.
II. A logical address (IP) is used for source-to-destination delivery across networks.
III. A port address identifies a specific process on a host computer.
IV. A specific address might be a URL for a web page.
V. A logical address has a global or universal scope.
Which of the above statements are correct?
(A) I, II, and III only
(B) I, II, III, and IV only
(C) I, II, III, IV, and V
(D) I, III, and V only

121. The term 'latency' in a data communication system refers to:
(A) The rate of data transfer.
(B) The number of bits in a signal.
(C) The delay in transmitting a message from source to destination.
(D) The capacity of the channel.

122. Which of the following is an example of an unguided transmission medium?
(A) Coaxial Cable
(B) Radio Waves
(C) Fiber-Optic Cable
(D) Twisted Pair

123. Which of the following is NOT a layer in the TCP/IP protocol suite?
(A) Application Layer
(B) Transport Layer
(C) Network Interface Layer
(D) Session Layer

124. The token passing access method is used in which of the following network topologies?
(A) Ring and Bus
(B) Bus and Star
(C) Mesh and Star
(D) Star only

125. Match List I (Technology) with List II (Description).
List I: (Technology)
P. FDMA
Q. TDMA
R. CDMA
S. CSMA
List II: (Description)
1. Divides access by time slots.
2. Spreads signal over the entire spectrum using a unique code.
3. A contention-based method where stations listen before talking.
4. Divides the channel spectrum into frequency bands.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-1, R-2, S-3
(C) P-4, Q-2, R-1, S-3
(D) P-2, Q-1, R-3, S-4

126. An IPv4 datagram has a header of 24 bytes. What is the value of the HLEN field?
(A) 6
(B) 5
(C) 24
(D) 96

127. Which protocol provides reliable, connection-oriented service at the transport layer?
(A) UDP
(B) IP
(C) ARP
(D) TCP

128. What is a key feature of SCTP that is not present in TCP or UDP?
(A) Multi-homing
(B) Connectionless communication
(C) 8-bit port numbers
(D) Use in the application layer

129. A trojan horse is a type of malware that:
(A) Replicates itself across the network.
(B) Encrypts user files and demands a ransom.

NODIA

(C) Disguises itself as legitimate software.
(D) Requires a host file to spread.

130. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : In symmetric-key cryptography, the same key is used for both encryption and decryption.
Reason (R) : The main challenge in symmetric-key cryptography is the secure distribution of the shared key.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

131. Consider the following technologies and protocols:
I. SMTP
II. GPRS
III. Firewall
IV. SaaS
V. ARP
Which of the above are directly related to Network Security or Mobile Technology?
(A) I, II, and III only
(B) II, III, and IV only
(C) II and III only
(D) I, IV, and V only

132. The architecture of GSM does not include which of the following?
(A) Network and Switching Subsystem (NSS)
(B) Virtual Private Network (VPN)
(C) Base Station Subsystem (BSS)
(D) Operation Support Subsystem (OSS)

133. Which cloud service model provides the highest level of control and management over the infrastructure for the user?
(A) SaaS
(B) PaaS
(C) IaaS
(D) FaaS

134. Which of the following correctly pairs a network device with its primary operating layer in the OSI model?
(A) Hub - Data Link Layer
(B) Switch - Network Layer
(C) Repeater - Transport Layer
(D) Router - Network Layer

135. Consider the following statements about IPv4 addressing:
I. Class A addresses have an 8-bit network prefix.
II. Class D addresses are reserved for multicasting.
III. The address 127.0.0.1 is a loopback address.
IV. Classless Inter-Domain Routing (CIDR) was introduced to mitigate the rapid exhaustion of IPv4 addresses.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and IV only
(C) I, II, III, and IV
(D) III and IV only

136. Which one of the following is a function of the Data Link Layer?
(A) Routing
(B) Flow control
(C) Congestion control
(D) Process-to-process delivery

137. Which of the following is a characteristic of a circuit-switched network?
(A) Data is sent in small packets.
(B) It is highly efficient for bursty data traffic.
(C) No resources are allocated until data is actually sent.
(D) A dedicated path is established for the duration of the connection.

138. The number of layers in the OSI reference model is:
(A) 7
(B) 4
(C) 5
(D) 6

139. A device connecting two or more networks and forwarding packets between them is called a:
(A) Hub
(B) Bridge

Click Here to Buy book on Amazon

(C) Router
(D) Repeater

140. Consider the following items related to network protocols and standards:
I. HDLC
II. CSMA/CA
III. SMTP
IV. FDMA
V. Virtual Private Network
Which of the above is NOT a multiple access method?
(A) I, III, and V
(B) II and IV
(C) I, II, and IV
(D) III and V

141. The function of fragmentation and reassembly of data is handled by which layer in the OSI model?
(A) Data Link Layer
(B) Transport Layer
(C) Session Layer
(D) Network Layer

142. The Uniform Resource Locator (URL) specifies which of the following?
(A) Only the IP address of the server.
(B) The protocol to be used, the domain name/ IP address, and the path of the resource.
(C) The physical MAC address of the destination server.
(D) The size and type of the file being requested.

143. To ensure a digital signature provides non-repudiation, the sender must encrypt the message digest with their:
(A) Public Key
(B) The receiver's Public Key
(C) Private Key
(D) A shared symmetric key

144. In the context of mobile computing, which component acts as a bridge between two different networks, such as a mobile network and the internet?
(A) Middleware
(B) Mobile IP
(C) GPRS
(D) Gateway

145. Arrange the following cloud computing models in decreasing order of user control over the underlying infrastructure:
1. Infrastructure as a Service (IaaS)
2. Platform as a Service (PaaS)
3. Software as a Service (SaaS)
Choose the correct sequence:
(A) 3, 2, 1
(B) 1, 2, 3
(C) 2, 1, 3
(D) 1, 3, 2

146. A private cloud is one where the cloud infrastructure is:
(A) Owned and managed by a third-party provider and shared by many organizations.
(B) Provisioned for exclusive use by a single organization.
(C) A composition of two or more distinct cloud infrastructures (private, community, or public).
(D) Shared by several organizations and supports a specific community that has shared concerns.

NODIA

147. A key benefit of virtualization is:
(A) Improved server utilization and energy efficiency.
(B) Increased hardware costs.
(C) Reduced software licensing options.
(D) Increased complexity in network management.

148. The term for the maximum amount of data that can be transmitted over a network link, often referred to as the link's payload capacity, is:
(A) ARP (Address Resolution Protocol)
(B) TTL (Time to Live)
(C) MTU (Maximum Transmission Unit)
(D) DNS (Domain Name System)

149. Which of the following is NOT a service provided by TCP?
(A) Flow Control
(B) Error Control
(C) Connection-oriented service
(D) Multicasting

150. Match List I (Security Term) with List II (Definition).
List I: (Security Term)

(A) It relies on a centralized base station.
(B) It is designed for high-speed, long-distance communication via satellite.
(C) It is a self-organizing, infrastructure-less network of mobile nodes.
(D) It uses physical cables to connect nodes.

170. Consider the following security mechanisms:
I. Digital Signature
II. Firewall
III. Secret-Key Algorithm
IV. Public-Key Algorithm
V. Steganography
Which of the above are security mechanisms related to Network Security?
(A) I, II, III, and IV only
(B) I, II, III, IV, and V
(C) I, III, and IV only
(D) II, IV, and V only

171. A network that connects devices within a person's immediate workspace is a:
(A) LAN
(B) MAN
(C) WAN
(D) PAN (Personal Area Network)

172. The term wireless network broadly refers to any computer network where:
(A) All devices are mobile.
(B) The interconnection between nodes is implemented without the use of wires.
(C) Data is transmitted only via satellite.
(D) The network topology is always a mesh.

173. What is the relationship between the OSI model and the TCP/IP suite?
(A) The OSI model is an implementation of the TCP/IP suite.
(B) The TCP/IP suite was designed based on the OSI model.
(C) The OSI model is a theoretical reference model, while TCP/IP is a practical, implemented protocol suite.
(D) They are two completely unrelated concepts in networking.

174. Which of the following is NOT a function of the transport layer?
(A) Process-to-process delivery
(B) Congestion control
(C) Error control
(D) Framing

175. Match List I (Protocol/Standard) with List II (Layer).
List I: (Protocol/Standard)
P. FTP
Q. TCP
R. IP
S. Ethernet
List II: (Layer)
1. Network Layer
2. Application Layer
3. Data Link Layer
4. Transport Layer
Choose the correct option from those given below:
(A) P-2, Q-4, R-1, S-3
(B) P-4, Q-2, R-1, S-3
(C) P-2, Q-1, R-4, S-3
(D) P-1, Q-4, R-2, S-3

176. Which of the following is an example of classless addressing?
(A) A Class B network.
(B) A network with address 10.0.0.0/8.
(C) A network with address 192.168.1.0 and mask 255.255.255.0.
(D) A Class D multicast address.

177. If a sender's TCP window size is 4000 bytes and the receiver acknowledges receipt of 1000 bytes, what can the sender do under the sliding window mechanism?
(A) Stop sending until the full 4000 bytes are acknowledged.
(B) Reduce the window size to 3000 bytes.
(C) Discard the connection.
(D) Send another 1000 bytes.

178. What does resolution in the context of DNS refer to?
(A) The process of mapping a domain name to an IP address.
(B) The screen resolution of the client computer.
(C) The negotiation of transmission speed.
(D) The fixing of transmission errors.

179. A software that is secretly installed on a computer to gather information about the user without their knowledge is best described as:

Downloaded from www.pdf.tube

Reason (R) : In PaaS, the cloud provider manages the underlying infrastructure, including servers, storage, and networking, as well as the operating system and middleware, allowing developers to focus solely on their applications.

In light of the above statements, choose the most appropriate answer from the options given below:

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

161. Consider the following communication scenarios:

I. A keyboard sending characters to a computer.
II. A telephone conversation between two people.
III. A walkie-talkie conversation between two security guards.
IV. A radio station broadcasting music.
V. A remote control changing a TV channel.

Which of the above are examples of Simplex communication?

(A) I, II, and III
(B) I, IV, and V
(C) II and III
(D) IV and V only

162. Which data encoding technique converts digital data into analog signals for transmission over an analog medium?

(A) Data Modulation
(B) Data Compression
(C) Data Encryption
(D) Data Framing

163. A switch operates at which layer of the OSI model?

(A) Physical Layer
(B) Network Layer
(C) Data Link Layer
(D) Transport Layer

164. What is the purpose of the polling multiple access method?

(A) To allow stations to transmit whenever they have data.
(B) To pass a special frame, the token, to authorize transmission.
(C) To divide the channel into time slots for stations.
(D) To have a master station control when other stations can transmit.

NODIA

165. Consider the following network devices:

I. A hub broadcasts incoming packets to all of its ports.
II. A switch forwards packets only to the specific port of the destination device.
III. A router connects different networks and makes routing decisions based on IP addresses.
IV. A bridge operates at the Network layer to connect two different LANs.

Which of the above statements are correct?

(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

166. Which field in the IPv4 header is used to prevent a packet from looping indefinitely in the network?

(A) Checksum
(B) Time to Live (TTL)
(C) Protocol
(D) Version

167. Indirect network layer delivery occurs when:

(A) The source and destination are on the same local network.
(B) The packet is delivered using a physical address only.
(C) The network is using circuit switching.
(D) The source and destination are on different networks and the packet must go through a router.

168. Which protocol is responsible for transferring web pages from a server to a browser?

(A) HTTP
(B) SMTP
(C) FTP
(D) TELNET

169. What is the main characteristic of a MANET (Mobile Adhoc Network)?

(D) LAN, MAN, and WAN services.

210. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).

Assertion (A) : The Internet of Things (IoT) has significant implications for data generation and network traffic.

Reason (R) : IoT involves a massive number of devices, many of which continuously collect and transmit small amounts of sensor data, leading to a large volume of traffic.

In light of the above statements, choose the most appropriate answer from the options given below:

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

211. Consider the following statements about TCP and UDP:

I. TCP guarantees ordered delivery of packets.
II. UDP is faster than TCP due to lower overhead.
III. TCP performs congestion control.
IV. UDP is suitable for applications like live video streaming where speed is more critical than reliability.

Which of the above statements are correct?

(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

212. Which mode of communication is used in a standard keyboard and monitor connection?

(A) Full-duplex
(B) Simplex
(C) Half-duplex
(D) Multiplex

213. A switch builds its forwarding table by examining which part of an incoming frame?

(A) Source IP address
(B) Destination IP address
(C) Source MAC address
(D) Destination Port number

214. Which layer of the OSI model is concerned with the syntax and semantics of the information transmitted?

(A) Application Layer
(B) Session Layer
(C) Transport Layer
(D) Presentation Layer

215. Consider the following statements about IPv6:

I. It uses a 128-bit address space.
II. The header includes a checksum field for error checking.
III. It supports autoconfiguration of addresses.
IV. Fragmentation is primarily handled by the source host, not routers.

Which of the above statements are correct?

(A) I and II only
(B) I, III, and IV only
(C) II and III only
(D) I, II, III, and IV

216. Which of the following is a primary characteristic of a Wide Area Network (WAN)?

(A) Confined to a single room or building.
(B) Spans a large geographical area, often connecting multiple LANs.
(C) High-speed network covering a city.
(D) Exclusively uses wireless technology.

217. When a mobile user moves from their home network to a different network, the Mobile IP protocol uses a Home Agent and a:

(A) Foreign Agent
(B) Mobile Agent
(C) Gateway Agent
(D) Roaming Agent

218. An SLA (Service Level Agreement) is crucial in cloud computing because it:

(A) Defines the cost of the service.
(B) Acts as a user manual for the service.
(C) Defines the performance, availability, and responsibilities guaranteed by the provider.
(D) Is a license to use the cloud software.

219. The process of converting analog signals into a digital format is known as:

(A) Modulation
(B) Demodulation
(C) Multiplexing
(D) Digitization (e.g., Pulse Code Modulation)

Click Here to Buy book on Amazon

(A) A firewall
(B) A virus
(C) Spyware (a type of malware)
(D) A VPN

180. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : GPRS was an important evolution for GSM networks.
Reason (R) : GPRS introduced packet-switched data services, enabling always-on connectivity and more efficient use of the radio network for data transmission compared to circuit-switched data.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

181. Consider the following statements about wireless LANs (WLANs):
I. They use radio waves to transmit data.
II. They are commonly used to provide network access in homes and offices.
III. The IEEE 802.11 family of standards governs most WLAN technologies.
IV. They are inherently more secure than wired LANs.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

182. Which of these is a core service provided by GSM?
(A) Telephony (voice calls)
(B) Video conferencing
(C) Cloud storage
(D) Platform as a Service

183. Which of the following is a key characteristic of the IoT?
(A) A small, fixed number of connected devices.
(B) Communication exclusively between humans.
(C) A focus on analog data transmission.
(D) A vast network of uniquely identifiable, interconnected devices.

184. In data communication, what is the effect of noise?
(A) It increases the speed of transmission.
(B) It can corrupt the signal, leading to errors in the received data.
(C) It enhances the signal strength.
(D) It is a necessary component for digital transmission.

185. Consider the following statements about network topologies:
I. In a mesh topology, a failure of one link does not affect the entire network.
II. A star topology's main vulnerability is the failure of the central hub.
III. In a ring topology, data typically circulates in one direction.
IV. A bus topology is easy to expand by adding new devices to the main cable.
Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

186. Which type of address is represented as 48 bits (6 bytes) and is often written in hexadecimal format?
(A) MAC Address
(B) IP Address (v4)
(C) Port Address
(D) URL

187. The Sliding Window protocol combines which two functions?
(A) Framing and Error Detection
(B) Addressing and Routing
(C) Multiplexing and Modulation
(D) Flow Control and Error Control

188. Which of the following is NOT a controlled access method?
(A) CSMA/CD
(B) Reservation
(C) Polling
(D) Token Passing

Downloaded from www.pdf.tube

189. Which part of an IPv4 address is used by routers for forwarding packets?
(A) The host portion
(B) The checksum field
(C) The network portion
(D) The protocol field

190. Consider the following network application protocols:
I. SMTP
II. FTP
III. TELNET
IV. POP
V. IMAP
Which of these protocols are part of the electronic mail architecture?
(A) I, II, and III
(B) I, IV, and V
(C) II, III, and IV
(D) All of the above

191. The primary purpose of using communication satellites is to:
(A) Provide very low-latency communication.
(B) Create small, ad-hoc networks.
(C) Connect devices within a single building.
(D) Overcome geographical barriers and provide wide area coverage.

192. Which mobile communication protocol allows a user to keep the same IP address while moving from one network to another?
(A) GPRS
(B) Mobile IP
(C) SMS
(D) GSM

193. A company develops a CRM application and hosts it on the cloud for its employees to use via a web browser. This is an example of what kind of cloud model?
(A) Public Cloud, SaaS
(B) Private Cloud, IaaS
(C) Private Cloud, SaaS
(D) Public Cloud, PaaS

194. A developer rents virtual servers, operating systems, and databases from a cloud provider to build and deploy a new web application. This is an example of:
(A) SaaS
(B) IaaS
(C) On-premise hosting
(D) PaaS

195. Arrange the following events in the typical resolution of a domain name (e.g., www.example.com) by a client:
1. The client's resolver queries a root DNS server.
2. The client's resolver contacts the authoritative DNS server for 'example.com'.
3. The client receives the IP address for 'www.example.com'.
4. The client's resolver checks its local cache.
Choose the most logical sequence, assuming the name is not in the cache:
(A) 4, 1, 2, 3
(B) 1, 2, 3, 4
(C) 4, 2, 1, 3
(D) 2, 1, 4, 3

196. The difference between simplex and half-duplex is that:
(A) Simplex is two-way; half-duplex is one-way.
(B) Simplex is one-way; half-duplex is two-way but not at the same time.
(C) Simplex is for analog signals; half-duplex is for digital signals.
(D) Simplex requires a physical medium; half-duplex is always wireless.

NODIA

197. Which error detection method involves arranging data into a table and calculating parity for each row and column?
(A) Two-Dimensional Parity Check
(B) Simple Parity Check
(C) Checksum
(D) Cyclic Redundancy Check

198. Which of the following is the most accurate statement regarding VLANs?
(A) VLANs are created by physically rewiring the network switches.
(B) All devices in a VLAN must be connected to the same physical switch.
(C) VLANs allow a network administrator to logically segment a network, improving security and managing broadcast traffic.
(D) VLANs increase the amount of broadcast traffic on a network.

199. What is the primary role of the Checksum field in an IPv4 datagram?
(A) To verify the integrity of the packet's payload.
(B) To determine the time-to-live for the packet.
(C) To identify the destination port number.
(D) To verify the integrity of the packet's header.

200. Match List I (Mobile Tech Term) with List II (Description).
List I: (Mobile Tech Term)
P. GSM
Q. MANET
R. GPRS
S. SMS
List II: (Description)
1. A packet-oriented data service on 2G/3G networks.
2. A text messaging service component.
3. A global standard for mobile communications using TDMA.
4. An infrastructure-less, self-configuring network.
Choose the correct option from those given below:
(A) P-1, Q-2, R-3, S-4
(B) P-3, Q-4, R-2, S-1
(C) P-3, Q-4, R-1, S-2
(D) P-2, Q-1, R-3, S-4

201. Which cloud model would be most appropriate for an organization that needs to rent raw computing power (virtual machines, storage) to host its custom legacy applications?
(A) SaaS
(B) IaaS
(C) PaaS
(D) FaaS

202. What is the primary function of a communication satellite in a network?
(A) To function as a repeater in the sky.
(B) To store large amounts of data in orbit.
(C) To process data using onboard computers.
(D) To manage IP address allocation globally.

203. Which of the following is a key service offered under the GSM standard?
(A) Infrastructure as a Service
(B) Root DNS resolution
(C) Short Message Service (SMS)
(D) Public-key encryption

204. The core idea of Cloud Storage is:
(A) Storing data on your personal computer's hard drive.
(B) A new type of physical flash drive.
(C) Storing data in a compressed format on a single server.
(D) Storing data on a network of data servers accessible through the internet.

205. If you see an IP address like 172.16.10.5/24, what does the /24 signify?
(A) The number of hosts allowed on the network.
(B) The address is part of a Class B network.
(C) The first 24 bits represent the network portion of the address.
(D) The packet has a time-to-live of 24 hops.

206. Which of the following is a function of the Network Layer?
(A) Logical addressing and routing
(B) Physical addressing and framing
(C) Process-to-process delivery
(D) Data encryption and formatting

207. Which statement correctly identifies an incorrect pair?
(A) TCP - Connection-oriented
(B) UDP - Connectionless
(C) ARP - Maps logical to physical address
(D) DNS - Maps physical to logical address

208. What is the fundamental difference between public-key and secret-key cryptography?
(A) Public-key uses one key, secret-key uses two keys.
(B) Public-key uses a pair of keys (public and private), while secret-key uses a single shared key.
(C) Public-key is for digital signatures, secret-key is for encryption.
(D) Public-key is faster than secret-key.

209. The service architecture of GSM is divided into which three main categories?
(A) IaaS, PaaS, and SaaS.
(B) Public, Private, and Hybrid services.
(C) Bearer services, Teleservices, and Supplementary services.

Click Here to Buy book on Amazon

P. Malware
Q. Cryptography
R. Steganography
S. Firewall

List II: (Definition)
1. Hiding the existence of data.
2. Software designed to disrupt, damage, or gain unauthorized access.
3. A network security device that monitors and filters traffic.
4. The art of writing or solving codes for secure communication.

Choose the correct option from those given below:
(A) P-1, Q-3, R-2, S-4
(B) P-2, Q-4, R-1, S-3
(C) P-2, Q-1, R-4, S-3
(D) P-3, Q-4, R-1, S-2

151. Consider the following statements about transmission errors:
I. Single-bit errors affect one bit in a data unit.
II. Burst errors affect two or more consecutive bits.
III. Error handling mechanisms can include error detection and error correction.
IV. Transmission media are perfectly immune to errors.

Which of the above statements are correct?
(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

152. Which network covers a geographical area like a city or a large campus?
(A) MAN
(B) LAN
(C) WAN
(D) PAN

153. In the TCP/IP protocol suite, which layer is responsible for the functionalities of the OSI model's Session, Presentation, and Application layers?
(A) Transport Layer
(B) Internet Layer
(C) Application Layer
(D) Host-to-Network Layer

154. What is a primary function of the HDLC protocol?
(A) To provide a standard for high-level programming.
(B) To define a bit-oriented protocol for communication over point-to-point and multipoint links.
(C) To manage IP address allocation.
(D) To encrypt email messages.

155. A router determines the path for a packet to its destination by consulting a:
(A) MAC table
(B) DNS cache
(C) Routing table
(D) ARP cache

156. The checksum in the UDP header provides:
(A) Guaranteed delivery
(B) Flow control
(C) Congestion avoidance
(D) Error detection for the header and data

157. Which of the following is a remote login protocol?
(A) TELNET
(B) FTP
(C) SMTP
(D) DNS

158. In CDMA, how are different users' signals separated?
(A) By assigning different time slots.
(B) By assigning a unique code to each user.
(C) By assigning different frequency bands.
(D) By using different physical cables.

159. A wireless network that is formed on the fly without any pre-existing infrastructure is called a:
(A) Cellular Network
(B) Wireless LAN (WLAN)
(C) Mobile Adhoc Network (MANET)
(D) Virtual Private Network (VPN)

160. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : PaaS (Platform as a Service) provides a higher level of abstraction than IaaS (Infrastructure as a Service).

NODIA

220. Consider the following items:
I. SaaS
II. PaaS
III. IaaS
IV. Public Cloud
V. Private Cloud
Which of the above represent cloud service models?
(A) I, II, and III only
(B) IV and V only
(C) I, II, III, IV, and V
(D) I and IV only

221. In a frequency division multiplexing (FDM) system, the channels are separated by unused frequency strips called:
(A) End bands
(B) Guard bands
(C) Channel gaps
(D) Signal dividers

222. Which protocol would be used to retrieve email from a mail server and allows management of mail folders on the server itself?
(A) IMAP
(B) SMTP
(C) POP3
(D) FTP

223. A VPN creates a secure tunnel across a public network. What is the most common technique used to secure this tunnel?
(A) Steganography
(B) Physical security
(C) MAC address filtering
(D) Encryption

224. Which of the following best describes Resource Management in the context of cloud computing?
(A) The process of hiring and managing IT staff.
(B) The efficient and dynamic allocation of computing resources (CPU, memory, storage) to users as needed.
(C) Managing the billing and subscription for cloud services.
(D) The physical organization of servers in a data center.

225. Match List I (Network Attack/Malware) with List II (Description).
List I: (Network Attack/Malware)
P. Virus
Q. Worm
R. Trojan Horse
S. Spyware
List II: (Description)
1. A self-replicating program that spreads across networks without a host file.
2. A program that collects user information without their consent.
3. A malicious program that attaches itself to another executable file.
4. A program that appears legitimate but performs a malicious function.
Choose the correct option from those given below:
(A) P-3, Q-1, R-4, S-2
(B) P-1, Q-3, R-4, S-2
(C) P-3, Q-4, R-1, S-2
(D) P-4, Q-1, R-3, S-2

226. What is the primary advantage of classless addressing (CIDR) over classful addressing?
(A) It is simpler to understand.
(B) It allows for more efficient and flexible allocation of IP addresses.
(C) It eliminates the need for routers.
(D) It only works with IPv6.

227. In the TCP header, which flag is used to initiate a connection?
(A) ACK
(B) FIN
(C) RST
(D) SYN

228. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The Data Link layer is responsible for node-to-node delivery of frames.
Reason (R) : The Data Link layer uses physical addresses (MAC addresses) to send frames from one node to the next node in the path.
In light of the above statements, choose the most appropriate answer from the options given below:
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.

Downloaded from www.pdf.tube

II. The TCP/IP suite is a 4-layer (or 5-layer) practical model.
III. The OSI model was developed by ISO as a framework for international standards.
IV. In practice, the TCP/IP model is more widely implemented than the OSI model.
V. Both models feature a Network/Internet layer for routing packets.

Which of the above statements are correct?
(A) I, II, and III only
(B) I, II, III, and V only
(C) II, III, and IV only
(D) I, II, III, IV, and V

Direction : Please read the following source carefully.

A data communication system is fundamentally composed of five key components: the message, sender, receiver, transmission medium, and a set of protocols. The efficacy of this system is heavily influenced by its communication mode. In simplex mode, communication is unidirectional, akin to a one-way street, where one device transmits and the other only receives. Half-duplex mode allows both devices to transmit and receive, but not simultaneously; communication alternates between the two parties. Full-duplex mode offers the most flexible communication, permitting simultaneous two-way data flow. The signals traversing the transmission medium can be either analog, characterized by continuous waves, or digital, represented by a sequence of discrete voltage pulses. The performance of this communication is measured by metrics such as bandwidth, which indicates the channel's capacity; throughput, which measures the actual rate of successful message delivery; and latency, the delay from signal transmission to reception. Real-world channels can be noisy, introducing errors, whereas an ideal noiseless channel, though theoretical, provides a benchmark for maximum data rate calculations. These fundamental concepts form the bedrock upon which all complex data communication architectures are built.

Based on the above source, attempt the following five questions.

251. Which communication mode allows for simultaneous two-way data transmission between devices?
(A) Simplex
(B) Full-duplex
(C) Half-duplex
(D) Unidirectional

252. According to the passage, what is the primary distinction between bandwidth and throughput?
(A) Bandwidth represents the theoretical capacity of a channel, while throughput is the actual successful data transfer rate.
(B) Bandwidth is a measure of delay, while throughput measures channel capacity.
(C) Bandwidth applies only to digital signals, whereas throughput applies only to analog signals.
(D) Bandwidth is measured in seconds, whereas throughput is measured in Hertz.

NODIA

253. A television broadcast system where the station sends signals to viewers' televisions without receiving any data back is a practical example of which communication mode?
(A) Half-duplex mode
(B) Full-duplex mode
(C) Simplex mode
(D) Duplex mode

254. The source describes a digital signal as being represented by:
(A) A set of continuous waves.
(B) The physical transmission medium itself.
(C) The protocols governing communication.
(D) A sequence of discrete voltage pulses.

255. What does the concept of latency in a data communication system refer to?
(A) The total capacity of the transmission channel.
(B) The rate at which messages are successfully delivered.
(C) The time delay between sending a signal and it being received.
(D) The presence of errors introduced into the signal.

Direction : Please read the following source carefully.

The transmission of data across a network involves intricate processes of encoding, modulation, and multiplexing. Data transmission can be digital, sending data bits directly, or analog, which involves converting digital data into an analog

unreliable.

Reason (R) : UDP does not provide acknowledgements, error correction, or retransmission of lost packets.

In light of the above statements, choose the most appropriate answer from the options given below:

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

241. Consider the following statements about Firewalls:

I. They can be implemented as hardware, software, or both.
II. They form a barrier between a secure internal network and an untrusted external network.
III. Packet-filtering firewalls make decisions based on source/destination IP address and port numbers.
IV. Firewalls can prevent all types of malware attacks.

Which of the above statements are correct?

(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

242. In the OSI model, which layer is responsible for managing dialogues between two computers?

(A) Session Layer
(B) Transport Layer
(C) Presentation Layer
(D) Application Layer

243. The cellular topology is a key concept in:

(A) Wired LANs
(B) Cloud computing
(C) Mobile communications
(D) Network security

244. Which of the following is NOT a valid reason to use UDP over TCP?

(A) The application requires low latency.
(B) The application can tolerate some data loss.
(C) The application involves multicasting or broadcasting to many hosts.
(D) The application requires guaranteed, in-order delivery of data.

245. Consider the following statements regarding the World Wide Web:

I. URL is used to uniquely identify a resource on the web.
II. DNS translates domain names into IP addresses.
III. HTTP is the protocol used to transfer web documents.
IV. HTML is the language used to create web pages.

Which of the above statements are correct?

(A) I and II only
(B) I, II, and III only
(C) III and IV only
(D) I, II, III, and IV

246. What type of addresses are used by the transport layer to identify the source and destination application processes?

(A) Physical addresses
(B) Port addresses
(C) Logical addresses
(D) Specific addresses

247. A virtual server is best described as:

(A) An emulated server instance running on a physical host.
(B) A physical server that is managed remotely.
(C) A backup server kept offline.
(D) A server with no operating system.

248. Digital signatures provide all of the following EXCEPT:

(A) Authentication
(B) Integrity
(C) Confidentiality
(D) Non-repudiation

249. Which of the following is an example of Baseband transmission?

(A) Cable Television
(B) AM/FM Radio
(C) Satellite communication
(D) Ethernet LAN

250. Consider the following statements about network architecture models:

I. The OSI model is a 7-layer reference model.

Click Here to Buy book on Amazon

(D) A is false but R is true.

229. Which of the following is NOT a valid IPv4 class in the classful addressing scheme?
(A) Class A
(B) Class C
(C) Class F
(D) Class E

230. A device that connects dissimilar networks by translating between different networking protocols is a:
(A) Gateway
(B) Hub
(C) Switch
(D) Repeater

231. The term mobile computing refers to:
(A) The use of cloud services.
(B) The architecture of the GSM network.
(C) A specific security protocol for wireless networks.
(D) A technology that allows transmission of data, voice, and video via a computer or any other wireless-enabled device without being connected to a fixed physical link.

232. What is a key difference between GSM and CDMA?
(A) GSM is for data, CDMA is for voice.
(B) GSM uses TDMA and FDD, while CDMA uses spread spectrum technology.
(C) GSM is an American standard, CDMA is a European standard.
(D) GSM requires a physical SIM card, while CDMA does not.

233. The concept of a Virtual Server is a direct application of which technology?
(A) Multiplexing
(B) Cryptography
(C) Virtualization
(D) Routing

234. In the context of IoT, what is a primary function of sensors?
(A) To process large amounts of data.
(B) To display information to the user.
(C) To provide network connectivity.
(D) To collect data from the physical environment.

235. Arrange the following networks in ascending order of their typical geographical size:
1. WAN
2. MAN
3. LAN

Choose the correct sequence:
(A) 1, 2, 3
(B) 3, 2, 1
(C) 2, 3, 1
(D) 3, 1, 2

236. The process by which DNS servers find the IP address for a given hostname is called:
(A) Allocation
(B) Resolution
(C) Registration
(D) Delegation

237. Which of the following provides the most basic cloud service, offering fundamental building blocks like virtual machines and storage?
(A) PaaS
(B) IaaS
(C) SaaS
(D) FaaS

NODIA

238. In a token ring network, what is the purpose of the token?
(A) It gives a station permission to transmit data.
(B) It carries the data from source to destination.
(C) It is used to acknowledge receipt of a frame.
(D) It identifies the address of the central hub.

239. What is the primary difference between flow control and congestion control in TCP?
(A) Flow control is end-to-end, while congestion control is network-wide.
(B) Flow control uses the sliding window, while congestion control uses checksums.
(C) Flow control prevents a fast sender from overwhelming a slow receiver; congestion control prevents the network itself from being overloaded.
(D) There is no difference; the terms are interchangeable.

240. Given below are two statements: one is labelled as Assertion (A) and the other is labelled as Reason (R).
Assertion (A) : The UDP protocol is considered

signal. This conversion is achieved through data encoding and modulation techniques. Encoding techniques, such as various line coding schemes, convert digital data to digital signals optimized for transmission. Modulation techniques, on the other hand, are used to impress digital or analog information onto a carrier signal, making it suitable for analog transmission. A key distinction in transmission strategy is between broadband and baseband. Baseband transmission uses the entire bandwidth of the medium to transmit a single signal, typically in a digital format. In contrast, broadband transmission divides the bandwidth of the medium into multiple, independent channels, allowing several signals to be transmitted simultaneously. This is made possible through multiplexing, a technique that combines multiple data streams into a single, composite signal for transmission over a shared medium. Common multiplexing methods enable efficient use of the transmission medium's capacity. Furthermore, the medium itself can introduce transmission errors, which necessitates the implementation of error handling mechanisms to ensure data integrity.

Based on the above source, attempt the following five questions.

256. Which transmission method, according to the passage, uses the entire bandwidth of a medium to carry a single signal?
(A) Broadband transmission
(B) Baseband transmission
(C) Multiplexed transmission
(D) Modulated transmission

257. What is the primary function of multiplexing in data communication?
(A) To convert digital data into analog signals.
(B) To detect and correct errors in transmission.
(C) To use the entire bandwidth for a single, high-speed signal.
(D) To combine multiple data streams into a single signal for a shared medium.

258. Based on the description, what is the key difference between data encoding and modulation?
(A) Encoding converts digital data to digital signals, while modulation impresses information onto a carrier for analog transmission.
(B) Encoding is used in broadband transmission, while modulation is used in baseband transmission.
(C) Encoding handles error correction, while modulation manages data flow.
(D) Encoding is a form of multiplexing, while modulation is a form of data transmission.

259. Broadband transmission's ability to carry multiple signals simultaneously is enabled by which of the following?
(A) Using the entire bandwidth for one signal.
(B) Sending data bits directly without a carrier signal.
(C) Dividing the medium's bandwidth into separate channels.
(D) Implementing error handling mechanisms.

260. The passage states that the conversion of digital data into an analog signal is an example of:
(A) Analog transmission.
(B) Digital transmission.
(C) Baseband transmission.
(D) Error handling.

Direction : Please read the following source carefully.

Modern computer networks are structured around a layered architecture, a design philosophy that divides the complex task of networking into smaller, more manageable layers. Each layer is designed to offer specific services to the layers above it, while abstracting the details of how those services are implemented. The two most prominent network models are the OSI (Open Systems Interconnection) Reference Model and the TCP/IP Protocol Suite. The OSI model is a conceptual framework with seven distinct layers, each with a defined function, serving as a standard for protocol development. In contrast, the TCP/IP suite, which is the practical foundation of the modern internet, is typically described with four or five layers. It is a more pragmatic model derived from implementation rather than a prescriptive standard. A crucial aspect of this layered approach is addressing. Different layers use different types of addresses to identify entities. For instance, the Physical layer deals with the raw bitstream, the Data Link layer uses physical (or MAC) addresses, the Network layer uses logical (or

NODIA

273. Which of the following is NOT a channelization protocol according to the passage?
(A) FDMA
(B) TDMA
(C) CDMA
(D) CSMA/CA

274. In a scenario where a central device systematically asks each secondary device if it has data to send, which multiple access protocol is being used?
(A) Polling
(B) Token passing
(C) CSMA/CD
(D) CDMA

275. The purpose of multiple access protocols is to:
(A) Encrypt data for secure transmission.
(B) Route packets between different networks.
(C) Coordinate access to a shared communication channel.
(D) Detect and correct errors in data frames.

Direction : Please read the following source carefully.

The Network Layer is responsible for host-to-host delivery of packets, potentially across multiple networks. Central to this is the Internet Protocol (IP), with IPv4 being the most historically widespread version. The IPv4 structure defines a 32-bit address, creating a finite address space. This space was traditionally managed using classful addressing, which divided addresses into fixed classes (A, B, C, etc.) with predefined network and host portions, leading to inefficient address allocation. To resolve this, classless addressing (CIDR) was introduced, allowing for variable-length subnet masking and more flexible allocation. An IPv4 packet, or datagram, may be too large for an underlying network's frame size, necessitating fragmentation, where the datagram is broken into smaller pieces. The IPv4 header includes a checksum field, used to check for corruption in the header itself. The newer IPv6 features a vastly larger 128-bit address space and a simplified packet format to improve efficiency. For a packet to be sent on a local network, the Network Layer's logical (IP) address must be mapped to the Data Link Layer's physical (MAC) address. This is the function of the Address Resolution Protocol (ARP). Routing algorithms are employed by routers to determine the best path for these datagrams to travel from the source to the final destination, enabling both direct and indirect network layer delivery.

Based on the above source, attempt the following five questions.

276. What was the primary motivation for introducing classless addressing (CIDR) to replace classful addressing?
(A) To address the inefficient allocation of IP addresses caused by fixed classes.
(B) To increase the size of the IPv4 header.
(C) To eliminate the need for the checksum field in the datagram.
(D) To enable the mapping of logical addresses to physical addresses.

277. The process of breaking a large datagram into smaller pieces to fit within a network's frame size is known as:
(A) Checksum
(B) Resolution
(C) Routing
(D) Fragmentation

278. What is the main function of the Address Resolution Protocol (ARP) as described in the passage?
(A) To determine the best path for a packet through a network.
(B) To map a logical IP address to a physical MAC address.
(C) To check the IPv4 header for errors.
(D) To fragment large datagrams into smaller packets.

279. Which feature of IPv4 is used to verify the integrity of the packet header?
(A) The fragmentation offset
(B) The classful address structure
(C) The checksum field
(D) The routing algorithm

280. Compared to IPv4, the most significant advantage of the IPv6 packet format mentioned in the source is its:
(A) Use of classful addressing.
(B) Requirement for fragmentation in all cases.
(C) Inclusion of a more complex checksum calculation.

Downloaded from www.pdf.tube

(C) To manage the transmission rate between sender and receiver.
(D) To allow a sender to transmit multiple frames before receiving an acknowledgment.

267. The Sliding Window Protocol is highlighted as an improvement over other methods because it:
(A) Only performs error detection without needing flow control.
(B) Combines flow control and error control, allowing multiple frames to be sent before acknowledgment.
(C) Is a bit-oriented protocol exclusively used for multipoint links.
(D) Prevents a fast sender from overwhelming a slow receiver by stopping all transmission.

268. According to the source, which of the following is a comprehensive, standardized protocol that performs framing, flow control, and error control at the Data Link Layer?
(A) High-Level Data Link Control (HDLC)
(B) Sliding Window Protocol
(C) Cyclic Redundancy Check (CRC)
(D) Checksum

269. If a fast computer is sending data to an older, slower printer, which Data Link Layer function is essential to prevent data loss?
(A) Framing
(B) Error detection
(C) Flow control
(D) Error correction

270. What is the logical sequence of actions in error control as implied by the passage?
(A) Correction is performed, then detection, followed by retransmission.
(B) Retransmission is always the first step, followed by detection.
(C) A checksum is calculated after the data is corrected to ensure accuracy.
(D) Detection mechanisms identify corruption, which then triggers correction or retransmission.

Direction : Please read the following source carefully.

In broadcast networks where a communication channel is shared by multiple stations, a mechanism is required to coordinate access and avoid collisions. This coordination is the responsibility of the Media Access Control (MAC) sublayer of the Data Link Layer, which employs multiple access protocols. These protocols can be broadly categorized into random access, controlled access, and channelization protocols. Among random access methods, CSMA/CD (Carrier Sense Multiple Access with Collision Detection) is famously used in traditional Ethernet. Stations listen before transmitting and can detect collisions, stopping and retrying after a random time. Its wireless counterpart, CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance), attempts to avoid collisions by using acknowledgment signals, as collision detection is difficult in wireless environments. Controlled access methods, such as polling and token passing, provide more orderly access. In polling, a central controller queries each station for data to send. In token passing, a special frame called a token circulates among stations, and only the station holding the token is allowed to transmit. Channelization protocols like FDMA (Frequency-Division Multiple Access), TDMA (Time-Division Multiple Access), and CDMA (Code-Division Multiple Access) divide the shared channel's bandwidth, time, or code space among different stations, allowing simultaneous, collision-free transmission.

Based on the above source, attempt the following five questions.

271. What is the fundamental difference between CSMA/CD and CSMA/CA?
(A) CSMA/CD is used for controlled access, while CSMA/CA is for random access.
(B) CSMA/CD is used in wireless networks, while CSMA/CA is used in wired Ethernet.
(C) CSMA/CD detects collisions, while CSMA/CA attempts to avoid them before they happen.
(D) CSMA/CD divides the channel by time, while CSMA/CA divides it by frequency.

272. A network protocol where a special frame must be possessed by a station before it can transmit data is an example of:
(A) Polling
(B) Token passing
(C) CSMA/CD
(D) FDMA

NODIA

protocols: Simple Mail Transfer Protocol (SMTP) is used for sending mail from a client to a server and for transferring mail between servers. For retrieving mail from a server, clients typically use either Post Office Protocol (POP), which generally downloads and deletes mail from the server, or Internet Message Access Protocol (IMAP), which allows users to manage mail directly on the server, synchronizing it across multiple clients. Other important protocols at this layer include TELNET, for remote terminal access, and FTP (File Transfer Protocol), for transferring files between systems.

Based on the above source, attempt the following five questions.

286. What is the primary function of the Domain Name Service (DNS)?
(A) To transfer files between different computer systems.
(B) To manage and synchronize emails directly on a server.
(C) To identify the location of resources on the World Wide Web using a URL.
(D) To translate human-readable domain names into machine-readable IP addresses.

287. Which protocol would a user's email client most likely use to send an outgoing email to their mail server?
(A) POP
(B) IMAP
(C) SMTP
(D) FTP

288. According to the passage, a key difference between POP and IMAP is that:
(A) POP is used for sending mail, while IMAP is for receiving.
(B) POP typically removes mail from the server after download, while IMAP manages it on the server.
(C) IMAP is an older protocol that has been replaced by POP.
(D) IMAP translates domain names, while POP transfers files.

289. The passage defines a Uniform Resource Locator (URL) as a means to:
(A) Perform the translation of names to addresses.
(B) Provide remote terminal access to a server.
(C) Send electronic mail between servers.
(D) Identify resources on the World Wide Web.

290. A user wants to access their email from their phone, laptop, and desktop computer, ensuring that actions like reading or deleting an email on one device are reflected on all others. Which mail retrieval protocol is best suited for this need?
(A) IMAP
(B) SMTP
(C) POP
(D) TELNET

Direction : Please read the following source carefully.

Network security is a critical discipline focused on protecting network integrity, accessibility, and the data transmitted across it from unauthorized access, misuse, or destruction. Threats come in many forms, including malwares, which are malicious software designed to disrupt operations or gain access to systems. To counter these threats, several security techniques are employed. Cryptography is the science of secure communication, using algorithms to transform plaintext data into an unreadable format (ciphertext) and back again, ensuring confidentiality. It is distinct from steganography, which is the art of hiding data within another file or message, concealing the very existence of the communication. Cryptographic algorithms are broadly classified into two types: secret-key (or symmetric) algorithms, where the same key is used for both encryption and decryption, and public-key (or asymmetric) algorithms, which use a pair of keys—a public key for encryption and a private key for decryption. Public-key cryptography also enables the creation of a digital signature, a mechanism used to verify the authenticity and integrity of a digital message. For securing communications over public networks like the internet, Virtual Private Networks (VPNs) can create an encrypted tunnel. Firewalls act as a barrier between a trusted internal network and an untrusted external network, controlling incoming and outgoing traffic based on a set of security rules.

Based on the above source, attempt the following five questions.

Click Here to Buy book on Amazon

(D) Vastly larger 128-bit address space.

Direction : Please read the following source carefully.

The Transport Layer provides logical communication between application processes running on different hosts. The three main protocols at this layer are TCP, UDP, and SCTP, each offering a different level of service. TCP (Transmission Control Protocol) is a connection-oriented protocol that provides reliable, ordered, and error-checked delivery of a stream of bytes. To achieve this reliability, it employs mechanisms for flow control, error control, and congestion control. Flow control, often using a sliding window, prevents a fast sender from overwhelming a slow receiver. Error control ensures data integrity through checksums and acknowledgments with retransmission. Congestion control mechanisms are vital for network stability, allowing TCP to reduce its transmission rate when it detects network congestion. In contrast, UDP (User Datagram Protocol) is a connectionless protocol that provides a very simple, unreliable service. It offers no flow control, error control (beyond an optional checksum), or congestion control, making it fast and suitable for applications where speed is more critical than reliability, such as streaming or DNS. SCTP (Stream Control Transmission Protocol) combines features of both TCP and UDP. It provides reliable, message-oriented transport like TCP but also supports multi-streaming and multi-homing, offering enhanced reliability and performance for certain applications. Both TCP and SCTP incorporate sophisticated congestion control to manage network traffic effectively.

Based on the above source, attempt the following five questions.

281. Which protocol is described as connectionless and providing a simple, unreliable service without flow or congestion control?
(A) UDP
(B) TCP
(C) SCTP
(D) ARP

282. What is the primary purpose of congestion control in TCP and SCTP?
(A) To ensure that data arrives in the correct order.
(B) To reduce the transmission rate in response to network traffic overload.
(C) To prevent a fast sender from overwhelming a slow receiver.
(D) To detect and correct bit errors within a packet.

283. According to the passage, TCP uses a sliding window primarily for which purpose?
(A) Congestion control
(B) Flow control
(C) Error control
(D) Connection setup

284. A key similarity between TCP and SCTP mentioned in the source is that both:
(A) Incorporate congestion control mechanisms.
(B) Are connectionless protocols.
(C) Provide an unreliable service.
(D) Lack any form of error control.

285. For a real-time online game where low latency is crucial and the occasional loss of a packet is acceptable, which transport protocol would be the most suitable choice based on the descriptions?
(A) TCP, because it guarantees delivery.
(B) SCTP, because it supports multi-homing.
(C) UDP, because it is fast and does not have the overhead of reliability mechanisms.
(D) TCP, because it provides effective flow control.

Direction : Please read the following source carefully.

The Application Layer provides services directly to the user and their applications. A cornerstone of this layer is the World Wide Web (WWW), a system for accessing and navigating interlinked documents. Resources on the WWW are identified by a Uniform Resource Locator (URL). To access these resources, a human-readable domain name (e.g., www.example.com) must be translated into a machine-readable IP address. This critical translation is performed by the Domain Name Service (DNS), a hierarchical and distributed naming system. DNS resolution can map names to addresses and, through reverse lookups, can also map addresses back to names. Another ubiquitous application layer service is Electronic Mail. The architecture of email involves several

NODIA

IP) addresses, and the Transport layer uses port addresses to identify specific processes on a host. Communication across networks is facilitated by switching techniques, which are methods for moving data through a series of intermediate nodes from source to destination.

Based on the above source, attempt the following five questions.

261. According to the passage, which type of address is primarily used at the Transport layer to identify specific processes?
(A) Physical address
(B) Logical address
(C) Specific address
(D) Port address

262. What is the primary advantage of a layered architecture in network design as described in the source?
(A) It increases the speed of data transmission by bypassing certain functions.
(B) It simplifies network design by dividing complex tasks into manageable, independent layers.
(C) It combines all networking tasks into a single, efficient protocol.
(D) It mandates the use of the seven-layer OSI model for all network implementations.

263. How does the passage contrast the OSI Reference Model with the TCP/IP Protocol Suite?
(A) OSI is a practical model used for the internet, while TCP/IP is a theoretical, seven-layer framework.
(B) OSI uses port addresses, whereas TCP/IP uses logical addresses.
(C) OSI is a conceptual, seven-layer model, while TCP/IP is a more pragmatic, implementation-derived suite.
(D) OSI is a switching technique, while TCP/IP is a layered architecture.

264. The use of logical addresses, such as IP addresses, is associated with which layer in a layered network model?
(A) Physical Layer
(B) Data Link Layer
(C) Transport Layer
(D) Network Layer

265. The passage defines switching techniques as methods for:
(A) Moving data through intermediate nodes from source to destination.
(B) Assigning physical and logical addresses to devices.
(C) Dividing networking tasks into a series of layers.
(D) Identifying specific processes running on a host computer.

Direction : Please read the following source carefully.

The lower layers of network models, particularly the Data Link Layer, are responsible for reliable node-to-node data delivery. A fundamental function at this level is framing, which involves encapsulating network layer datagrams into frames, adding headers and trailers that contain control information. This process is essential for providing a structure for the bitstream. Another critical responsibility is error control, which encompasses both error detection and correction. Techniques for error detection, such as checksums or cyclic redundancy checks (CRC), are used to identify if data has been corrupted during transmission. Upon detection, mechanisms for correction or retransmission are initiated. Closely related is flow control, which manages the rate of data transmission between a fast sender and a slow receiver to prevent the receiver from being overwhelmed. The Sliding Window Protocol is a sophisticated method that combines flow control with error control. It allows the sender to transmit multiple frames before needing an acknowledgment, improving efficiency over stop-and-wait protocols. High-Level Data Link Control (HDLC) is a standardized, bit-oriented protocol that implements many of these functions, including framing, flow control, and error control, providing a comprehensive solution for point-to-point and multipoint link management.

Based on the above source, attempt the following five questions.

266. What is the primary purpose of framing as described in the passage?
(A) To detect and correct errors in the data.
(B) To encapsulate network layer datagrams with headers and trailers for structured transmission.

Downloaded from www.pdf.tube

291. What is the fundamental difference between cryptography and steganography as described in the passage?
(A) Cryptography uses a public key, while steganography uses a secret key.
(B) Cryptography is a type of malware, while steganography is a defense against it.
(C) Cryptography makes data unreadable, while steganography hides the existence of the data.
(D) Cryptography is used for firewalls, while steganography is used for VPNs.

292. Which security mechanism is described as acting like a barrier that controls traffic between a trusted and an untrusted network?
(A) A Virtual Private Network (VPN)
(B) A Firewall
(C) A Digital Signature
(D) Steganography

293. In a public-key algorithm, how are the keys typically used for sending a confidential message?
(A) The sender's private key is used for encryption, and the receiver's private key is used for decryption.
(B) The same secret key is used for both encryption and decryption.
(C) The receiver's private key is used for encryption, and the receiver's public key is used for decryption.
(D) The receiver's public key is used for encryption, and the receiver's private key is used for decryption.

294. A digital signature, which is enabled by public-key cryptography, is primarily used to:
(A) Encrypt the entire content of a message for confidentiality.
(B) Verify the authenticity and integrity of a message.
(C) Hide a message within an image file.
(D) Create a secure, encrypted tunnel over the internet.

295. The term malwares in the context of network security refers to:
(A) The science of secure communication using algorithms.
(B) A set of security rules applied by a firewall.
(C) Malicious software designed to disrupt systems or steal information.
(D) The pair of keys used in asymmetric cryptography.

Direction : Please read the following source carefully.

NODIA

The landscape of communication is increasingly shaped by mobile technology, cloud computing, and the Internet of Things (IoT). Mobile technologies like GSM (Global System for Mobile Communications) and CDMA (Code Division Multiple Access) provide the underlying infrastructure for wireless cellular communication. The architecture of GSM, for instance, defines the services and structure for mobile computing, which is further enabled by middleware and gateways that bridge mobile devices with other networks and services. A key protocol, Mobile IP, allows mobile devices to maintain a constant IP address while moving between different networks. On a larger scale, communication satellites facilitate global connectivity. This mobile ecosystem is intrinsically linked with cloud computing, which offers on-demand access to a shared pool of configurable computing resources. Cloud services are categorized into models like SaaS (Software as a Service), PaaS (Platform as a Service), and IaaS (Infrastructure as a Service), and can be deployed in public or private clouds. A core technology enabling the cloud is virtualization, which allows for the creation of virtual servers and abstracts resources like cloud storage and database storage. The relationship between a cloud provider and a user is governed by a Service Level Agreement (SLA). These technologies converge in the Internet of Things (IoT), a network of interconnected physical devices that collect and exchange data, fundamentally extending the reach of the internet into the physical world.

Based on the above source, attempt the following five questions.

296. According to the passage, what is the primary role of virtualization in cloud computing?
(A) To define the terms of service in a Service Level Agreement.
(B) To provide the underlying infrastructure for GSM and CDMA networks.
(C) To categorize cloud services into SaaS, PaaS, and IaaS models.

single signal for transmission over a shared medium. This process allows for more efficient use of available bandwidth and reduces the need for multiple physical links. At the receiving end, the signals are separated through demultiplexing, ensuring each reach their intended destination correctly across the network.

10. (A) Cyclic Redundancy Check (CRC)
Cyclic Redundancy Check is a powerful mathematical technique used for detecting errors in transmitted data blocks. It involves calculating a short, fixed-length check value based on the contents of the frame. This value is appended to the message, allowing the receiver to perform a similar calculation to verify that no bits were corrupted during the transmission through the communication channel.

11. (D) I, II, III, and IV
Baseband transmission uses the entire bandwidth for one signal and is typically digital, serving short distances. In contrast, broadband transmission divides the available bandwidth into multiple frequency bands to carry several signals simultaneously using analog techniques. Both methods are essential concepts in networking, used for different scales and types of data delivery across various physical media, making all four statements accurate.

12. (B) Actual rate at which data is successfully transferred.
Throughput is a critical performance metric that measures the actual amount of data successfully transferred from one point to another within a given timeframe. Unlike theoretical bandwidth, which represents the maximum potential capacity of a channel, throughput accounts for overhead, errors, and network congestion. It provides a more realistic assessment of how well a communication system is performing during actual use.

13. (C) A mesh topology is the least reliable due to single points of failure.
Mesh topology provides high reliability because every device is connected to multiple others, creating redundant paths for data. If one link fails, the system can reroute traffic through alternative connections, ensuring continuous operation. This makes the claim that it is the least reliable due to single points of failure incorrect, as it actually lacks such vulnerabilities compared to other topologies.

14. (D) Wide Area Network (WAN)
A Wide Area Network is specifically designed to span large geographical distances, such as those between different cities, countries, or even continents. Companies use WANs to interconnect their local networks across diverse locations, often utilizing leased lines, satellite links, or public networks. This makes it the most appropriate choice for connecting corporate offices that are not within the same metropolitan area.

15. (A) Layered Area Network
Networks are classified by their geographical span, such as Local Area Networks for small areas and Wide Area Networks for larger regions. Metropolitan Area Networks cover city-sized domains. The term Layered Area Network is not a standard classification used in computer networking to describe the physical or geographical scope of a network, making it the incorrect option among the standard network types.

16. (B) Layered architecture
Layered architecture is a design principle that organizes network communication tasks into a vertical stack of distinct layers. Each layer performs a specific set of functions and interacts only with the layers directly above and below it. This approach simplifies network design, makes the system more modular, and allows different protocols to be developed and updated independently without affecting the entire system.

17. (B) III, IV, and V only
The OSI model is a theoretical framework consisting of seven layers that define network communications. It includes the physical layer as its foundation. While the OSI model is modular and aids in design, the TCP/IP suite is the protocol stack actually implemented on the internet. Therefore, statements about OSI's seven layers, modularity, and the physical layer's position are correct, while others are incorrect.

Click Here to Buy book on Amazon

ANSWERS

1. (B) Sender
A data communication system relies on several vital elements to function effectively. The sender is the specific device or entity responsible for initiating the transmission by creating and sending the data or message. Other critical components include the receiver, the transmission medium, the message itself, and the protocols that govern the exchange of information across the communication network.

2. (A) Full-duplex
Full-duplex communication allows data to be transmitted and received simultaneously in both directions over a single channel. This mode is highly efficient as it avoids waiting for the channel to be clear, unlike simpler modes. Common examples of this technology include modern telephone networks and high-speed internet connections that support concurrent upstream and downstream data traffic in real-time environments.

3. (C) They are discrete, time-varying signals.
Digital signals are characterized by having discrete values, typically represented as binary zeros and ones, rather than a continuous range. These signals vary over time to represent information in a structured format. They offer better resistance to noise and interference compared to analog signals, making them ideal for modern computer networks and long-distance digital data transmission systems used globally.

4. (D) Half-duplex
A walkie-talkie operates in a mode where communication can occur in both directions, but only one party can transmit at a given time. This requires users to take turns speaking, as the device cannot send and receive signals simultaneously. This specific type of transmission is a fundamental concept used in various radio and early networking environments for communication.

5. (A) P-4, Q-2, R-1, S-3
Bandwidth is the range of frequencies a channel can pass, whereas throughput represents the actual measured rate of data transfer. Latency refers to the time delay for a message to arrive at its destination. Simplex describes a unidirectional communication channel. These definitions are fundamental concepts that characterize the performance and directionality of various data communication systems and modern computer networks.

6. (A) Both Statement I and Statement II are true.
Nyquist's formula indicates that even in a noiseless channel, the bit rate is limited by bandwidth and signal levels. Shannon's theorem states that a noisy channel's capacity depends on both bandwidth and the signal-to-noise ratio. Both these principles are foundational to communication theory and accurately describe the theoretical limits of data transmission in different channel conditions for network systems.

7. (D) To encode digital data onto an analog carrier signal.
Modulation is the process of modifying a high-frequency carrier signal to represent information. Its primary purpose in data transmission is to encode digital data onto an analog carrier signal, allowing it to be transmitted over mediums that are naturally analog, such as radio waves or telephone lines. This technique is essential for long-distance communication and efficient bandwidth usage in networks.

8. (C) A is true but R is false.
In data communication, the transmission medium is the physical path through which a message travels between the sender and the receiver. This medium can be tangible, such as twisted-pair wires or coaxial cables. However, it can also be intangible, such as the atmosphere or space used for wireless signals. Therefore, the assertion is correct, but the reasoning is false.

9. (C) Multiplexing
Multiplexing is a technique used to combine multiple signals from different sources into a

Downloaded from www.pdf.tube

for modern use because it transmits all data, including passwords, in unencrypted text.

64. (D) FTP is for file transfer, while TELNET is for remote login.
File Transfer Protocol and TELNET are both foundational application layer protocols, but they serve different primary purposes. FTP is designed specifically for the efficient transfer of files between a client and a server, providing features for directory management and file manipulation. In contrast, TELNET is focused on providing a remote terminal session, allowing a user to execute commands on a distant host machine.

65. (B) I, II, and III only
Electronic mail systems utilize several specialized protocols. The Simple Mail Transfer Protocol is a push protocol for sending messages. The Post Office Protocol is a pull protocol for retrieving them. The Internet Message Access Protocol provides more advanced features by allowing users to manage their emails directly on the server. FTP is used for general file transfers and is not part of mail architecture.

66. (B) Physical layer technologies for cellular networks.
GSM and CDMA are two fundamental physical layer technologies used in cellular networks to manage how multiple users share the available radio frequency spectrum. GSM uses a combination of time and frequency division to separate users, while CDMA employs spread spectrum technology with unique codes. These competing standards defined the second and third generations of mobile telecommunications before the global transition to LTE.

67. (D) Base Station Subsystem (BSS)
The Base Station Subsystem is a critical component of the GSM network architecture responsible for handling the radio traffic between mobile devices and the core network. It consists of the Base Transceiver Stations, which communicate directly with mobile phones, and the Base Station Controllers, which manage the radio resources and handle handoffs between different cells within their specific geographic service area.

68. (A) It provides services to applications beyond what is available from the operating system.
Middleware in mobile computing acts as an intermediary layer that provides common services and capabilities to applications that are not natively offered by the operating system. It simplifies the development process by handling complex tasks such as data management, communication between distributed systems, and session management. This allows developers to focus on application logic while ensuring compatibility across diverse mobile environments.

69. (C) Allow a mobile device to maintain its IP address while moving between different networks.
Mobile IP is a specialized communication protocol that allows a mobile device to move from one network to another while maintaining its permanent home IP address. This ensures that active sessions and connections are not dropped when the device changes its point of attachment. It works by using home and foreign agents to tunnel data to the device's current location seamlessly.

NODIA

70. (C) I, II, IV, and V
Mobile and wireless communication systems rely on diverse technologies. Satellites relay signals over vast distances, and cellular topologies divide areas into cells. Wireless LANs follow standard protocols like IEEE 802.11, and GPRS adds packet-switched data to mobile networks. While MANETs are wireless, they are infrastructure-less, meaning they do not require fixed hardware. Consequently, the other four mentioned statements are technically correct.

71. (D) Connect a mobile network to other networks, translating protocols if necessary.
In the context of mobile computing, a gateway serves as a vital bridge between the mobile network and other external networks, such as the internet. Its primary function is to route traffic between these different systems and perform any necessary protocol translations. This ensures that data from mobile devices can be correctly understood and delivered to servers on different types of networks.

72. (B) Short Message Service

18. (A) Transmitting raw bits over a communication channel.
The physical layer is the lowest level in the OSI model and is primarily concerned with the transmission and reception of raw bitstreams over a physical medium. It defines the electrical, mechanical, and functional specifications for the physical link. This includes details such as voltage levels, timing of signal changes, physical connectors, and the representation of binary ones and zeros as signals.

19. (C) Port address
In network communication, several types of addresses are used to ensure data reaches its destination. While physical and logical addresses identify devices on a network, a port address is used to identify a specific process or application running on a host. This allows the transport layer to deliver data to the correct software entity among many processes active on the computer.

20. (B) P-3, Q-4, R-2, S-1
Network communication layers perform specific roles for data delivery. The application layer provides user interfaces and services, while the transport layer handles process-to-process delivery. The network layer is responsible for logical addressing and routing across systems. Finally, the data link layer manages framing and physical addressing to ensure data reaches the correct adjacent node on the physical link, completing the logical mapping.

21. (C) Physical address
A Media Access Control address is a unique identifier assigned to a network interface controller for use as a network address in communications within a network segment. It is considered a physical address because it is typically hardcoded into the device hardware. MAC addresses are used at the data link layer to ensure that data is delivered to the correct physical device.

22. (B) Circuit switching
Circuit switching is a networking method that establishes a dedicated, continuous physical path between two communicating nodes for the entire duration of their session. Once the connection is established, the resources are reserved exclusively for those users, ensuring a constant bit rate and minimal delay. This technique is traditionally used in public switched telephone networks for voice communication between two parties.

23. (C) A is true but R is false.
The TCP/IP protocol suite is the actual implemented standard used for global internet communications. Although the OSI model is a highly valuable theoretical framework for understanding networking layers, it was developed after TCP/IP and is not the primary suite used for practical implementation. Thus, the assertion regarding TCP/IP status is true, but the reasoning about OSI practicality for implementation is false.

24. (A) Encapsulation
Encapsulation is the process where each layer in a networking model adds its own control information, usually in the form of headers and trailers, to the data received from the layer above. This information is essential for the corresponding layer at the receiving end to correctly process and deliver the data. As the data moves down the stack, it becomes increasingly nested.

25. (C) Data Link Layer
The data link layer is responsible for taking the packets received from the network layer and organizing them into manageable units called frames. This layer handles the physical addressing and ensures reliable transmission across the physical link by performing error detection and flow control. By creating frames, the data link layer allows the physical layer to transmit the data as bits.

26. (B) I, II, and IV only
Error control is a vital aspect of reliable data transmission, involving both detection and correction of corrupted bits. It is primarily managed at the data link and transport layers to ensure data integrity. While parity checks are used for basic detection, they are not a correction method. Therefore, the statements regarding detection, correction, and layer placement are accurate for standard network protocols.

27. (D) Flow control
The sliding window protocol is a sophisticated mechanism used for flow control in network communications. It allows a sender to transmit

Click Here to Buy book on Amazon

(D) To allow for the creation of virtual servers and the abstraction of resources.

297. Which of the following allows a mobile device to keep the same IP address while moving between networks?
(A) Mobile IP
(B) GSM
(C) CDMA
(D) SaaS

298. If a company subscribes to a service that provides a complete software application over the internet, which cloud computing model are they using?
(A) IaaS (Infrastructure as a Service)
(B) SaaS (Software as a Service)
(C) PaaS (Platform as a Service)
(D) Private Cloud

299. The passage describes the Internet of Things (IoT) as:
(A) A specific mobile communication protocol like GSM.
(B) A network of interconnected physical devices that collect and exchange data.
(C) A document governing the relationship between a cloud provider and user.
(D) A technology for creating virtual servers in a public cloud.

300. What is the purpose of a Service Level Agreement (SLA) in the context of cloud computing?
(A) To govern the relationship and define service terms between a provider and a user.
(B) To allow mobile devices to maintain a constant IP address.
(C) To abstract computing resources like storage and databases.
(D) To bridge mobile devices with other networks using middleware.

NODIA

exactly 32 bits long, allowing for a theoretical maximum of approximately 4.3 billion unique addresses. They are typically expressed in dotted-decimal notation, such as 192.168.1.1, for easier human reading and management.

37. (C) The address space is divided into five classes: A, B, C, D, and E.
Classful addressing was the original method used to divide the IPv4 address space into five distinct classes based on the first few bits of the address. Classes A, B, and C were used for unicast communication in networks of different sizes, while Class D was reserved for multicasting and Class E for experimental purposes. This system has largely been replaced by CIDR.

38. (B) Fragmentation
Fragmentation is the process used in IP networking to break down large datagrams into smaller pieces so they can pass through network segments with a smaller Maximum Transmission Unit. Each fragment is treated as an independent packet and includes information in its header that allows the destination host to reassemble them back into the original datagram once all pieces have arrived.

39. (D) To provide more addresses and overcome IPv4 address exhaustion.
The primary motivation for developing Internet Protocol version 6 was the rapid exhaustion of available IPv4 addresses due to the explosive growth of the internet and connected devices. IPv6 uses a 128-bit address space, which provides a virtually inexhaustible supply of unique addresses. This ensure that every device can have its own global IP, simplifying network management and supporting future expansion.

40. (B) III, IV, and V only
IPv4 uses 32-bit addresses, whereas IPv6 uses 128-bit addresses to provide a larger address space. Checksums are not mandatory in the IPv6 header to improve processing speed, and classful addressing is strictly an IPv4 concept. Additionally, in IPv6, fragmentation is performed only by the source host rather than by routers, which makes the specific statements about IPv6 header and behavior incorrect.

41. (C) ARP
The Address Resolution Protocol is a critical networking protocol used to map a known logical IP address to its corresponding physical MAC address on a local area network. When a device needs to send data to another host on the same network, it uses ARP to find the physical address required for the data link layer to create and deliver frames correctly.

42. (A) Direct delivery
In network layer communication, direct delivery occurs when the source and destination hosts are located on the same physical network. In this scenario, the sender can deliver the packet directly to the receiver without the need for an intermediate router. The sender uses the receiver's physical address, obtained through protocols like ARP, to transmit the data frame across the shared medium.

43. (D) Sliding Window
Routing algorithms are used by network devices to determine the best path for data packets to travel across interconnected networks. Distance Vector, Link State, and Path Vector are all well-known classes of routing protocols. However, the sliding window is a mechanism used at the transport and data link layers for flow and error control, rather than for determining network routing paths.

44. (B) UDP
The User Datagram Protocol is a core member of the Internet Protocol suite that provides a connection less and unreliable service at the transport layer. UDP does not establish a formal connection before sending data and does not provide mechanisms for acknowledgments, retransmissions, or flow control. This makes it faster and more suitable for applications like live streaming and online gaming where speed is essential.

45. (B) 2, 1, 3
Establishing a TCP connection requires a specific sequence known as the three-way handshake. First, the client initiates the process by sending a SYN segment to the server. The server then responds with a SYN-ACK segment to acknowledge the request and synchronize its own parameters. Finally, the client sends an

Downloaded from www.pdf.tube

ACK segment back to the server to finalize the connection, ensuring both parties are ready.

46. (D) A sliding window
TCP uses a sliding window mechanism to implement flow control between a sender and a receiver. This allows the receiver to specify how much data it is currently able to accept, preventing the sender from overwhelming its available buffer space. The window size is dynamically adjusted based on the receiver's capacity, ensuring an efficient and balanced transfer of data across the connection.

47. (A) Preventing the network from being overwhelmed by excessive traffic.
In the context of TCP, congestion control is a mechanism designed to prevent the entire network from becoming overloaded with excessive traffic. It monitors the network's state and reduces the rate at which the sender transmits data if signs of congestion, such as packet loss or high delays, are detected. This differs from flow control, which focuses only on the receiver's buffer capacity.

48. (B) SCTP
Stream Control Transmission Protocol is a transport layer protocol that combines features of both TCP and UDP while adding its own advanced capabilities. Like TCP, it is connection-oriented and provides reliable delivery. However, SCTP also supports multi-homing, which allows a single connection to use multiple IP addresses, and multi-streaming, which prevents head-of-line blocking by allowing independent streams within one association.

49. (B) For error detection in the header.
The checksum field in the IPv4 datagram header is used exclusively for error detection within the header itself. It ensures that the critical information needed for routing, such as the source and destination IP addresses and the time-to-live value, has not been corrupted during transit. Routers verify this checksum at each hop, and if an error is detected, the packet is discarded.

50. (A) P-4, Q-1, R-2, S-3
Different protocols serve distinct roles within the networking stack. TCP provides connection-oriented, reliable transmission, while UDP offers a faster, connectionless datagram service. The Address Resolution Protocol maps logical IP addresses to physical MAC addresses. Finally, the Internet Protocol handles logical addressing and routing to ensure packets are delivered across various interconnected networks from the source to the final destination in a structured manner.

51. (C) Uniform Resource Locator
A Uniform Resource Locator is a standard way of identifying the location of a resource, such as a web page, image, or document, on the World Wide Web. It specifies the protocol used to access the resource, the domain name or IP address of the server, and the specific path to the file on that server. URLs are fundamental for web navigation.

52. (D) Resolve human-readable domain names into IP addresses.
The Domain Name System serves as the phonebook of the internet by translating human-readable domain names, such as www.example.com, into the numeric IP addresses that computers use to identify each other. This resolution process allows users to access websites using memorable names while the underlying network protocols handle the actual data routing using the required IP addresses.

53. (C) Mapping names to addresses and addresses to names in the system.
DNS resolution is a versatile system that provides both forward and reverse lookups. Forward resolution involves mapping human-readable hostnames to their corresponding numeric IP addresses, which is necessary for web browsing and other services. Reverse resolution allows the system to map an IP address back to a hostname. This dual capability ensures comprehensive name and address management across the global internet.

54. (B) SMTP
The Simple Mail Transfer Protocol is the standard protocol used in electronic mail architecture for sending messages from a client to an email server. It is also used by servers to relay mail to other servers across the internet. SMTP is a push protocol, meaning it is designed to move mail from the source towards the destination mail server.

NODIA

55. (D) Firewall
Cybersecurity involves protecting systems from various types of malicious software. Viruses, worms, and Trojan horses are all examples of malware designed to cause harm or gain unauthorized access to computer systems. A firewall, however, is a security tool or system designed to monitor and control incoming and outgoing network traffic based on predetermined rules, acting as a defense mechanism rather than a threat.

56. (A) To hide the message's existence within another file or medium.
Steganography is the art and science of hiding the very existence of a message or information within another medium, such as an image, audio file, or digital document. Unlike encryption, which makes a message unreadable, steganography aims to prevent an observer from even suspecting that a secret message is being communicated, providing an additional layer of security and privacy.

57. (C) Symmetric cryptography
Secret-key cryptography is also known as symmetric cryptography because it uses a single, identical key for both the encryption of plaintext and the decryption of ciphertext. Both the sender and the receiver must share this secret key in advance. While efficient for processing large amounts of data, the main challenge of this method lies in the secure distribution and management of the shared key.

58. (A) It uses two keys: a public key for encryption and a private key for decryption.
Public-key algorithms, also known as asymmetric cryptography, utilize a pair of mathematically related keys. The public key is made available to everyone and is used to encrypt data, while the private key is kept secret by the owner and is used for decryption. This approach solves the key distribution problem found in symmetric systems and enables secure communication without sharing secret keys.

59. (D) To provide authentication, non-repudiation, and integrity.
A digital signature is a cryptographic mechanism used to verify the authenticity and integrity of a digital document or message. It ensures that the sender is who they claim to be and that the content has not been altered during transmission. Additionally, it provides non-repudiation, meaning the sender cannot later deny having sent the message, which is essential for secure electronic transactions.

60. (A) Both A and R are true and R is the correct explanation of A.
A firewall is a critical security component that establishes a barrier between an internal trusted network and external untrusted networks like the internet. It operates by analyzing network traffic packets and comparing them against established security criteria. Based on these rules, the firewall either allows or blocks the traffic, effectively protecting the internal system from unauthorized access and potential external cyber threats.

61. (A) Create a secure, encrypted connection over a less secure network, like the internet.
A Virtual Private Network allows users to establish a secure and encrypted connection to another network over a public infrastructure like the internet. By creating a private tunnel for data, a VPN ensures that sensitive information remains confidential and protected from eavesdropping. It is widely used by businesses to provide remote workers with secure access to internal company resources from any location.

62. (B) IMAP
The Internet Message Access Protocol is a sophisticated protocol used for retrieving and managing electronic mail from a remote server. Unlike simpler protocols that often download and delete messages, IMAP allows users to view and organize their emails directly on the server. This enables seamless synchronization across multiple devices, as all changes to folders and message statuses are reflected globally.

63. (C) Providing a command-line interface to a remote computer.
TELNET is a long-standing network protocol used to provide a command-line interface for communicating with a remote computer. It allows a user to log into a distant host and execute commands as if they were physically present at that machine. While useful for administration and testing, TELNET is considered insecure

Click Here to Buy book on Amazon

multiple data packets before requiring an acknowledgment from the receiver. By dynamically adjusting the number of packets that can be in transit, it prevents a fast sender from overwhelming a slow receiver, thereby optimizing the use of available network resources.

28. (B) CSMA/CD
Carrier Sense Multiple Access with Collision Detection is a media access control method used in early Ethernet networks. It requires a station to listen to the channel before transmitting. If a collision occurs while sending, the station detects the overlap, stops transmission immediately, and waits a random amount of time before trying again. This process helps manage traffic on shared network mediums.

29. (B) II and V only
Networking devices operate at different layers of the OSI model. Hubs and repeaters function at the physical layer, dealing with raw signals. Routers operate at the network layer, managing logical addresses. Switches and bridges are designed to operate primarily at the data link layer, as they use physical MAC addresses to filter and forward data frames between different segments of a network.

30. (A) Both A and R are true and R is the correct explanation of A.
Routers are essential networking devices that operate at the network layer of the OSI model. Their primary role is to determine the best path for packets to travel between different networks. By using logical IP addresses, they make intelligent forwarding decisions, ensuring that data reaches its intended destination across complex, interconnected networks, which confirms the relationship between their layer and core function.

31. (A) CSMA/CA
Carrier Sense Multiple Access with Collision Avoidance is a multiple access method primarily used in wireless networks, such as those governed by the IEEE 802.11 standards. Since wireless devices often cannot detect collisions directly while transmitting, they use CA to minimize the chance of overlaps. This involves listening to the channel and using handshake signals to reserve the medium before sending data.

32. (B) To logically segment a LAN into different broadcast domains.
A Virtual Local Area Network is a technology that allows network administrators to logically segment a single physical LAN into multiple distinct broadcast domains. This segmentation is done through software and switch configuration rather than physical rewiring. VLANs improve network performance by reducing unnecessary broadcast traffic and enhance security by isolating sensitive departments or devices from the rest of the network.

33. (B) CDMA
Code Division Multiple Access is a channelization-based multiple access protocol that allows multiple users to share the same frequency band simultaneously. Instead of dividing the channel by time or frequency, CDMA assigns a unique mathematical code to each transmitter. This spreads the signal over the entire available spectrum, allowing the receiver to use the same code to extract the original data.

NODIA

34. (A) Data Link Layer
High-Level Data Link Control is a bit-oriented, synchronous data link layer protocol developed by ISO. It is used to provide reliable communication over point-to-point and multipoint links. HDLC handles framing, error control, and flow control, ensuring that data units are delivered accurately across a physical connection. It has served as the basis for many other important data link layer protocols.

35. (D) I, II, III, and IV
Multiple access methods regulate how devices share a communication channel. Polling involves a master device, while token passing uses a special frame for permission. Reservation allows stations to book future time slots. Contention-based methods like CSMA/CD require stations to listen before transmitting. All these techniques are standard approaches used to manage traffic and prevent collisions in shared medium environments, making the statements correct.

36. (D) 32 bits
An Internet Protocol version 4 address is a numeric identifier assigned to each device participating in a computer network that uses the IP for communication. These addresses are

Short Message Service is a standard communication protocol used for the exchange of short text messages between mobile devices. Originally developed as part of the GSM standards, SMS allows users to send and receive messages of up to 160 characters. It is one of the most widely used mobile services globally due to its simplicity, reliability, and compatibility across different networks and devices.

73. (C) Infrastructure-less and self-configuring.
A Mobile Adhoc Network is a decentralized type of wireless network that does not rely on any pre-existing infrastructure, such as base stations or access points. Instead, each node in a MANET participates in routing by forwarding data for other nodes. These networks are self-configuring and highly dynamic, as nodes can join or leave at any time, making them ideal for emergency situations.

74. (D) Determine the geographical position of a mobile device.
Wireless Geolocation Systems are technologies used to determine the exact or approximate geographical location of a mobile device. These systems can utilize various methods, including GPS satellite signals, cellular tower triangulation, or Wi-Fi network analysis. Geolocation is essential for many modern applications, such as navigation services, location-based advertising, emergency response tracking, and various logistics and asset management solutions used in mobile computing.

75. (C) P-3, Q-2, R-1
Cloud computing offers various service models based on user needs. Software as a Service provides applications over the internet, while Platform as a Service offers tools for developers to build applications without managing infrastructure. Infrastructure as a Service provides fundamental resources like virtual servers and storage. These models allow organizations to choose the level of control and management they require for their operations.

76. (B) Public Cloud
A public cloud is a cloud deployment model where the computing resources, such as servers and storage, are owned and operated by a third-party cloud service provider and delivered over the internet. These services are made available to the general public or large industry groups. Public clouds offer high scalability and cost-effectiveness, as users only pay for the specific resources they consume.

77. (A) Virtualization
Virtualization is a foundational technology in cloud computing that allows for the creation of multiple virtual instances of a physical resource, such as a server or operating system. By using a hypervisor, a single physical computer can run multiple independent operating systems simultaneously. This ensures efficient hardware utilization, provides isolation between different users, and enables the rapid scaling and flexibility characteristic of cloud environments.

NODIA

78. (C) Infrastructure as a Service
Infrastructure as a Service is a cloud computing model that provides fundamental computing, network, and storage resources on demand over the internet. Instead of purchasing and maintaining physical hardware, users can rent virtualized resources to host their applications and data. IaaS offers the highest level of flexibility and control among cloud models, allowing users to manage their own operating systems and software.

79. (D) A cloud service provider and a consumer.
A Service Level Agreement is a formal contract between a cloud service provider and a consumer that specifies the expected level of service. It defines critical performance metrics such as uptime, availability, and response times, as well as the responsibilities and penalties if those standards are not met. The SLA is essential for establishing trust and ensuring that business requirements are consistently fulfilled.

80. (A) Both A and R are true and R is the correct explanation of A.
Virtualization is a core technology that enables cloud computing by abstracting physical hardware into multiple virtual instances. This allows a single physical server to host many independent users and applications simultaneously. By pooling resources and providing elasticity, virtualization ensures that computing power is used efficiently and can scale rapidly to meet demand, which is

users to share the same frequency by transmitting in their assigned slots. This efficient use of the spectrum was a key factor in the global success of the second-generation mobile networks.

99. (C) SaaS
Software as a Service is a cloud computing model where an entire software application is provided to the customer over the internet. The application runs on the provider's cloud infrastructure, and the user accesses it via a web browser or client interface. In this model, the provider handles all maintenance, updates, and underlying infrastructure, allowing the user to focus solely on using the software.

100. (C) P-4, Q-2, R-3, S-1
Networking and security involve various specialized technologies. Virtual LANs are a type of LAN technology used to segment networks logically. Classful addressing is a specific historical method of allocating IPv4 addresses. Mobile IP is a communication protocol designed for mobile devices. Digital signatures are a fundamental tool in network security, used to provide authentication and ensure the integrity of transmitted digital messages.

101. (C) Operating System
A standard data communication system consists of five essential components: the message, the sender, the receiver, the transmission medium, and the protocol. The message is the information being sent, the sender initiates the transmission, the receiver is the destination, the medium is the physical path, and protocols are the rules. While important for computing, an operating system is not a fundamental component.

102. (B) BitRate = 2 * B * log2(L)
The Nyquist Bit Rate formula defines the theoretical maximum bit rate for a noiseless channel. It states that the bit rate is equal to twice the bandwidth multiplied by the base-2 logarithm of the number of signal levels used. This formula demonstrates that increasing either the bandwidth or the number of discrete signal levels will result in a higher data transmission capacity for the channel.

103. (D) Using analog signals and FDM to carry multiple channels.
Broadband transmission is characterized by its ability to carry multiple independent channels of information over a single physical medium simultaneously. This is achieved by using analog signaling and Frequency Division Multiplexing, where the total available bandwidth is divided into different frequency bands. This technology is commonly used in cable television and high-speed internet services to deliver various services over one connection.

104. (A) A break in the main cable can disable the entire network.
In a bus topology, all devices are connected to a single, central backbone cable. While this layout is simple and cost-effective to install, it has a significant disadvantage: the main cable represents a single point of failure. If there is a break or fault anywhere in this backbone, the entire network will be unable to communicate, as the shared transmission path is severed.

105. (C) LAN
A Local Area Network is a type of computer network that is restricted to a small geographical area, such as a single building, an office, or a university campus. LANs are characterized by high data transfer speeds and are typically owned and managed by the organization that uses them. They allow for the efficient sharing of resources like printers, files, and internet access.

106. (A) Presentation Layer
The presentation layer, which is the sixth layer of the OSI model, is responsible for managing the syntax and semantics of the information exchanged between two systems. Its primary functions include data translation, encryption for security, and compression to improve transmission efficiency. By handling these tasks, the presentation layer ensures that data sent from the application layer of one system is readable.

107. (D) MAC address
A Media Access Control address is a unique 48-bit identifier assigned to a network interface controller. It serves as a physical address that uniquely identifies a device on a local area network. MAC addresses are used at the data link layer to ensure that data frames are delivered to the correct hardware device on the physical

Downloaded from www.pdf.tube

based on the network's structure. This allows packets to be routed from a source to a destination regardless of the specific physical media used in between.

90. (A) I, II, III, IV, and V
Each layer of the OSI model has specific responsibilities to ensure successful communication. The data link layer handles framing, the network layer manages routing, and the transport layer controls congestion. The session layer is responsible for synchronization, while the presentation layer handles data encryption and formatting. Together, these seven layers provide a comprehensive framework for how data is processed and transmitted across network systems.

91. (D) Polling
Polling is a controlled access method used in networks where a central controller, often called the primary station, manages communication. The primary station invites other secondary stations to transmit data one by one. This centralized approach prevents collisions on the shared medium by ensuring that only the station explicitly selected by the controller is allowed to send its data at any given time.

92. (B) Networks with a small number of hosts.
In the original classful IPv4 addressing scheme, Class C addresses were specifically designed for networks with a relatively small number of devices. Each Class C network could support a maximum of 254 unique hosts. This allowed for a large number of individual networks to be created, making it suitable for smaller organizations or departments that did not require a vast address space.

93. (C) SCTP supports multi-streaming, while TCP does not.
A key difference between the Stream Control Transmission Protocol and TCP is that SCTP supports multi-streaming. This feature allows a single connection to carry multiple independent streams of data simultaneously. In TCP, a loss in one part of the data stream can block all subsequent data, but in SCTP, other streams can continue to flow, effectively eliminating the head-of-line blocking problem.

94. (D) Name translation
Name translation, commonly performed by the Domain Name System, is the process of mapping a human-readable domain name to its corresponding numeric IP address. This is also referred to as forward mapping. It allows users to access internet resources using memorable names like www.example.com while enabling the network's underlying protocols to handle the actual routing of data using the required IP addresses.

95. (A) 2, 4, 1, 3
During the encapsulation process at the sender's side, data travels down the OSI model. It begins at the application layer where the user interacts with the system. Then, it moves to the transport layer for segmenting, followed by the network layer for routing information. Finally, the data reaches the data link layer, where it is framed for transmission over the physical medium for the network.

96. (B) It is designed to download emails to the local client and then delete them from the server.
Post Office Protocol is an older email retrieval protocol that was designed based on the assumption that users would download their messages to a single local device. In its default operation, POP retrieves all new emails from the server, stores them on the client machine, and then removes them from the server. This makes it less suitable for users who access email from multiple devices.

NODIA

97. (B) B's public key
In public-key cryptography, to ensure that only the intended recipient can read a message, the sender must use the recipient's public key for encryption. Because only the recipient possesses the corresponding private key, they are the only party capable of decrypting the ciphertext. This ensures confidentiality and allows secure communication between parties who have never shared a secret key in advance.

98. (A) GSM
The Global System for Mobile Communications is a widely used digital cellular standard that is primarily based on Time Division Multiple Access. GSM divides each radio frequency channel into several time slots, allowing multiple

it. By overcoming signal attenuation, repeaters allow data to travel over longer physical distances than would otherwise be possible on a single cable.

117. (B) The practice of dividing a geographical service area into a grid of cells, each served by a base station.
A cellular topology is a fundamental design principle for mobile networks where a large geographical area is divided into smaller, hexagonal or circular regions called cells. Each cell is served by its own base station and uses a specific set of frequencies. This layout allows for frequency reuse in non-adjacent cells, significantly increasing the overall capacity and coverage of the mobile communication system.

118. (A) IoT focuses on connecting everyday physical devices to the internet to collect and share data.
The Internet of Things is a transformative concept that involves embedding sensors, software, and connectivity into everyday physical objects, from household appliances to industrial machinery. By connecting these devices to the internet, they can collect and exchange real-time data, enabling smarter automation, improved efficiency, and new types of services that were not possible with traditional, non-connected physical objects or systems.

119. (C) Software as a Service (SaaS)
Software as a Service is a cloud computing model where users access a fully functional application over the internet, typically through a web browser. Services like web-based email are classic examples of SaaS, as the provider manages the entire software application, its updates, and the underlying server infrastructure. The user is only responsible for using the service without any infrastructure management.

120. (C) I, II, III, IV, and V
Network addressing occurs at multiple levels to ensure proper data delivery. Physical addresses handle node-to-node transfers, while logical IP addresses manage source-to-destination routing across networks. Port addresses identify specific applications on a host, and specific addresses like URLs identify web resources. Logical addresses are designed to have a global scope, allowing for universal identification and communication across the vast, interconnected global internet system.

121. (C) The delay in transmitting a message from source to destination.
Latency is a key performance metric in data communication that refers to the total time delay experienced by a message as it travels from the source to its destination. It is composed of several factors, including propagation delay, processing delay at nodes, and queuing delay. Low latency is critical for real-time applications such as video conferencing, online gaming, and high-frequency financial trading.

122. (B) Radio Waves
Radio waves are an example of an unguided transmission medium, also known as wireless media. Unlike guided media like cables, unguided media transmit electromagnetic signals through the air or space without a physical conductor to direct them. Radio waves are widely used for various types of communication, including broadcasting, cellular networks, and wireless LANs, as they can travel over long distances.

123. (D) Session Layer
The TCP/IP protocol suite is organized into a four-layer or five-layer model that differs from the seven-layer OSI model. While it includes the application, transport, and internet layers, it does not have a distinct session or presentation layer. In the TCP/IP model, the functions of the OSI session and presentation layers are handled directly within the application layer protocols themselves.

124. (A) Ring and Bus
Token passing is a controlled access method used in networks with shared media, most notably in ring and bus topologies. In this scheme, a special frame called a token circulates among the nodes. A station can only transmit data when it possesses the token, ensuring that only one device communicates at a time. This prevents collisions and provides a predictable and efficient way.

125. (B) P-4, Q-1, R-2, S-3
Multiple access techniques define how users share the communication spectrum. FDMA assigns separate frequency bands, while TDMA uses specific time slots for each user. CDMA

Click Here to Buy book on Amazon

medium, regardless of the higher-level logical addressing.

108. (B) Both A and R are true but R is not the correct explanation of A.
Flow control is an essential mechanism used to regulate the speed of data transmission between a sender and a receiver. It prevents a fast sender from overwhelming a receiver that has limited processing or buffer capacity. The sliding window protocol is a standard technique used to implement this control. Both the assertion and the reason are true, though they describe separate related networking concepts.

109. (C) A centralized, high-capacity network that connects other smaller, independent networks.
A backbone network is a high-speed, high-capacity central infrastructure designed to interconnect various smaller networks, such as LANs or individual buildings within a large organization. It acts as the primary conduit for traffic traveling between these different network segments. Backbones are critical for ensuring efficient long-distance communication and maintaining the overall performance and scalability of a large-scale enterprise or internet network.

110. (A) A variable-length prefix or netmask.
In Classless Inter-Domain Routing, which replaced the older classful addressing system, the network portion of an IPv4 address is not fixed. Instead, it is determined by a variable-length prefix, often represented using a slash notation or a subnet mask. This allows for much more flexible and efficient allocation of IP addresses, as organizations can be assigned blocks that match their specific host requirements.

111. (D) It simplifies and speeds up packet processing by routers.
One of the key design improvements in IPv6 is the introduction of a fixed-size base header. Unlike the IPv4 header, which has a variable length due to optional fields, the IPv6 header is always 40 bytes. This consistency allows routers to process packets much more efficiently in hardware, reducing latency and increasing overall network throughput across the modern, high-speed global internet backbone.

112. (B) Indirect
In network layer communication, indirect delivery occurs when the source and destination hosts are located on different physical networks. In this case, the packet cannot be delivered directly and must instead be sent to an intermediate router. The router then uses its routing table to forward the packet towards the destination network, potentially passing through multiple hops before final delivery.

113. (C) A virus requires a host program to replicate, while a worm is a standalone program that can propagate itself.
Viruses and worms are both types of malicious software, but they differ in how they replicate and spread. A virus is a piece of code that must attach itself to an existing legitimate program or file to function and propagate. In contrast, a worm is a standalone program that can replicate itself and spread automatically across network connections without any user intervention.

NODIA

114. (D) Handoff
In cellular networks, a handoff is the process of transferring an active call or data session from one base station to another as the mobile user moves between different cells. A successful handoff must occur seamlessly and quickly enough that the user does not experience any interruption or loss in service. This capability is essential for providing continuous connectivity in mobile communication systems.

115. (D) I, II, III, and IV
Cloud deployment models define how cloud services are made available to users. Public clouds are open to everyone, while private clouds are dedicated to a single organization. Hybrid clouds combine different models to provide flexibility. Community clouds are shared by specific groups with common goals. All these models serve different organizational needs for security, control, and scalability in a cloud-based computing environment.

116. (B) Repeater
A repeater is a physical layer device used to extend the reach of a network. Its primary function is to receive an incoming signal that has been weakened by distance and regenerate it to its original strength and shape before retransmitting

or a large university campus. It is often used to connect multiple LANs within the same metropolitan region, providing high-speed connectivity over a city-wide scale.

153. (C) Application Layer
In the TCP/IP protocol suite, the application layer is a broad category that encompasses the functions performed by the application, presentation, and session layers of the seven-layer OSI model. It handles all high-level protocols used by software for network communication, including data formatting, encryption, and dialogue management. This consolidated approach reflects the practical design of the protocols used on the modern internet.

154. (B) To define a bit-oriented protocol for communication over point-to-point and multipoint links.
The primary function of the High-Level Data Link Control protocol is to provide a standardized, bit-oriented framework for reliable communication over various types of data links. HDLC manages how data frames are structured, ensuring they are correctly delimited and delivered across point-to-point and multipoint connections. It includes mechanisms for error detection and flow control, which are essential for maintaining data link integrity.

155. (C) Routing table
A router determines the best path for an incoming packet by consulting its internal routing table. This table contains a list of known networks and the corresponding next hop or interface to reach them. By comparing the packet's destination IP address with the entries in the routing table, the router makes an intelligent decision on how to forward the data towards its final destination.

156. (D) Error detection for the header and data
The checksum field in the UDP header is an optional feature used to provide error detection for both the header and the data payload. While UDP is a connectionless and unreliable protocol, the checksum allows the receiver to identify if the packet has been corrupted during transmission. If an error is detected, the receiver typically discards the packet, ensuring basic data integrity.

157. (A) TELNET
TELNET is a long-standing application layer protocol specifically designed for remote login. it allows a user on one computer to establish a virtual terminal session on a distant host machine. Once logged in, the user can execute commands as if they were physically present at the remote system's console. While historically important, its use has largely been replaced by more secure protocols like SSH.

158. (B) By assigning a unique code to each user.
In Code Division Multiple Access, different users' signals are separated by assigning a unique mathematical code to each transmitter. All users share the same frequency band simultaneously, but their data is multiplied by their specific code before transmission. This spreads the signal across the spectrum, and the receiver uses the same unique code to extract the original data from the combined signals.

159. (C) Mobile Adhoc Network (MANET)
A Mobile Adhoc Network is a type of wireless network that is formed spontaneously by mobile devices without the need for any pre-existing infrastructure, such as base stations or central controllers. Each node in a MANET acts as both a host and a router, forwarding data for other nodes to enable communication. These networks are self-configuring and highly flexible for dynamic environments.

160. (A) Both A and R are true and R is the correct explanation of A.
Platform as a Service provides a higher level of abstraction than Infrastructure as a Service. This is because the cloud provider manages not just the hardware and networking, but also the operating system and middleware. This allows developers to focus entirely on writing and deploying their applications without worrying about the underlying environment. Thus, the reason correctly explains the assertion regarding cloud model abstraction.

161. (B) I, IV, and V
Simplex communication is a unidirectional mode where data travels in only one direction. A keyboard sending characters to a computer is a classic example, as the data only flows from the device to the system. Radio broadcasting and television remote controls also use simplex

Downloaded from www.pdf.tube

employs spread spectrum technology with unique codes to separate signals. CSMA is a contention-based protocol where devices listen to the channel before transmitting. These methods are essential for managing traffic and optimizing the use of shared resources in networks.

126. (A) 6
The header length field in an IPv4 datagram specifies the length of the header in 4-byte increments. This field is 4 bits long, meaning its value must be multiplied by 4 to get the actual byte count. For a header of 24 bytes, the HLEN value would be 6, as six times four equals twenty-four. This allows for variable header sizes.

127. (D) TCP
The Transmission Control Protocol is a fundamental transport layer protocol that provides a connection-oriented and highly reliable service. Before sending data, TCP establishes a formal connection between the sender and receiver through a handshake process. It ensures reliability by using sequence numbers, acknowledgments, and retransmissions for any lost or corrupted segments, making it ideal for applications that require perfect data integrity.

128. (A) Multi-homing
Multi-homing is a unique feature of the Stream Control Transmission Protocol that is not found in the more traditional TCP or UDP protocols. It allows a single SCTP association to use multiple IP addresses for the same connection. This provides built-in redundancy, as the connection can automatically failover to an alternative IP path if one network link becomes unavailable or fails.

129. (C) Disguises itself as legitimate software.
A Trojan horse is a type of malicious software that achieves its goal by deceiving users about its true intent. It typically appears as a useful or harmless program, such as a utility or a game, but contains hidden malicious code. Once executed, it can perform various harmful actions, such as stealing data or providing a backdoor for attackers, without replicating itself.

130. (B) Both A and R are true but R is not the correct explanation of A.
Symmetric-key cryptography uses a single shared key for both the encryption and decryption processes. This makes the system efficient but introduces the significant challenge of securely sharing that key between the communicating parties. While both the assertion and the reason are true, they highlight different aspects of the technology rather than the reason providing a direct logical explanation for the assertion.

131. (C) II and III only
Among the provided options, GPRS and firewalls are directly related to mobile technology and network security, respectively. GPRS is a packet-switched data service for mobile networks, while firewalls are essential security devices used to filter network traffic. SMTP and ARP are general networking protocols for mail and address resolution, and SaaS is a cloud service model, making them less specific to those fields.

132. (B) Virtual Private Network (VPN)
The architecture of the Global System for Mobile Communications is composed of several major subsystems, including the Base Station Subsystem, the Network and Switching Subsystem, and the Operation Support Subsystem. These components work together to provide mobile connectivity and management. A Virtual Private Network, however, is a security technology used to create secure tunnels over networks and is not part of architecture.

133. (C) IaaS
Among the primary cloud service models, Infrastructure as a Service offers the highest level of control to the user. In the IaaS model, the provider offers the raw virtualized computing resources, but the user is responsible for managing the operating system, middleware, and applications. This flexibility makes it ideal for organizations that need to host custom or legacy systems with specific requirements.

134. (D) Router - Network Layer
Networking devices are designed to operate at specific layers of the OSI model based on the information they process. Hubs and repeaters function at the physical layer, while switches operate at the data link layer using MAC addresses. Routers are designed for the network layer, as they use logical IP addresses to make

NODIA

a fundamental requirement of modern, scalable cloud services.

81. (B) A network of physical objects using sensors to exchange data.
The Internet of Things refers to the vast network of physical objects—such as appliances, vehicles, and industrial equipment—that are embedded with sensors, software, and other technologies. These devices are connected to the internet, allowing them to collect, share, and exchange data without direct human intervention. The goal of IoT is to enable smarter decision-making and automation in various aspects of life.

82. (A) A software-based emulation of a physical server.
In cloud computing, a virtual server is a software-based representation of a physical server created through virtualization technology. It functions exactly like a real server, with its own operating system and applications, but it shares the underlying physical hardware with other virtual servers. This approach allows cloud providers to maximize resource efficiency and provide users with flexible, on-demand computing power.

83. (D) Data consistency, availability, and security
Database storage in a cloud environment presents unique challenges that providers must address to ensure reliable service. Maintaining data consistency across distributed systems is critical, as is ensuring high availability so that users can access their information at any time. Furthermore, robust security measures are essential to protect sensitive data from unauthorized access or loss, which are primary concerns for any cloud-based database.

84. (A) Packet-switched data service
General Packet Radio Service was an important evolution in mobile technology that introduced packet-switched data services to existing GSM networks. Unlike traditional circuit-switched connections used for voice calls, GPRS allows data to be transmitted in small packets, enabling more efficient use of network resources. This technology provided always-on connectivity for mobile internet access, paving the way for modern mobile data communication.

85. (B) I, II, and IV only
Network security encompasses several practices and tools to protect data. Cryptography secures communication, while steganography hides its very existence. Firewalls are versatile defenses that can be implemented in hardware or software to filter traffic. However, digital signatures rely on public-key cryptography rather than symmetric-key methods. Thus, the statements regarding cryptography, steganography, and firewalls are the ones that are accurate in this context.

86. (B) Full-duplex
A standard telephone conversation is a classic example of full-duplex communication. In this mode, both parties can transmit and receive voice data simultaneously over the same connection. This allows for natural, real-time interaction without the need for users to wait for the other person to finish speaking before they can reply, making it highly efficient for two-way human communication.

87. (D) Atmosphere
Transmission media are classified as either guided or unguided. Guided media use physical structures to direct signals, such as twisted-pair cables, coaxial cables, and fiber-optic cables. Unguided media, also known as wireless media, transmit signals through the air or space without a physical conductor. The atmosphere is the primary medium for unguided transmission, carrying radio waves, microwaves, and infrared signals.

88. (A) Mesh
Mesh topology is characterized by its high level of reliability and redundancy. In a full mesh network, every device is connected directly to every other device, ensuring that multiple paths exist for data transmission. If one link fails, the network can automatically reroute traffic through alternative paths. This robustness makes mesh topology ideal for critical networks where continuous connectivity and fault tolerance are required.

89. (C) Logical address
An Internet Protocol address is a logical address used at the network layer to uniquely identify a device across different interconnected networks. Unlike physical addresses, which are tied to specific hardware, logical addresses are assigned

Click Here to Buy book on Amazon

proves that the message indeed came from the specific sender and has not been altered.

144. (D) Gateway
In mobile computing, a gateway acts as a critical interface between different types of networks. Its primary role is to connect a mobile network to another network, such as the public internet, and perform any necessary protocol translations. This ensures that data from mobile devices can be correctly routed and understood by servers on external networks, enabling seamless global communication for mobile users.

145. (B) 1, 2, 3
Cloud computing models vary in the level of control provided to the user. Infrastructure as a Service offers the most control, including management of virtual servers and storage. Platform as a Service provides a middle ground for developers. Software as a Service offers the least control, as the provider manages the entire application and its underlying infrastructure. This hierarchy is standard in cloud service models.

146. (B) Provisioned for exclusive use by a single organization.
A private cloud is a cloud deployment model where the infrastructure is dedicated solely to one organization. It can be hosted internally on the company's own data center or managed by a third party. Private clouds offer the highest level of security and control, making them ideal for businesses with sensitive data or specific regulatory requirements that necessitate an isolated computing environment.

147. (A) Improved server utilization and energy efficiency.
A key benefit of virtualization is the significant improvement in server utilization and energy efficiency. By allowing multiple virtual machines to run on a single physical server, organizations can consolidate their hardware, reducing the total number of physical machines required. This leads to lower power consumption, reduced cooling costs, and more efficient use of available computing resources in data centers.

148. (C) MTU (Maximum Transmission Unit)
The Maximum Transmission Unit is a network setting that defines the largest size of a data unit, such as a packet or frame, that can be transmitted over a specific network link. It determines the maximum payload capacity of the link. If a packet exceeds the MTU of a network segment, it must be fragmented into smaller pieces before it can be sent.

149. (D) Multicasting
The Transmission Control Protocol provides several essential services for reliable data transmission, including flow control, error control, and connection-oriented communication. However, TCP is designed for point-to-point communication between two hosts and does not support multicasting. Multicasting, which involves sending a single stream of data to multiple recipients simultaneously, is typically handled by UDP or specialized network layer protocols.

150. (B) P-2, Q-4, R-1, S-3
Security terms define different aspects of system protection and threats. Malware refers to malicious software like viruses and worms. Cryptography is the practice of securing communications through codes. Steganography is a technique used to hide data within other files. A firewall is a network security device that monitors traffic. Together, these concepts form the foundation of modern cybersecurity and the protection of digital information.

151. (B) I, II, and III only
Data transmission errors can occur in different forms, such as single-bit errors or burst errors affecting multiple bits. Error handling mechanisms, including detection and correction, are essential for ensuring data integrity. However, no transmission medium is perfectly immune to errors, as noise, interference, and signal attenuation are inherent in all physical paths. Therefore, only the first three statements regarding error types and handling are correct.

152. (A) MAN
A Metropolitan Area Network is a computer network that interconnects users with computer resources in a geographic area larger than a local area network but smaller than a wide area network. Typically, a MAN spans an entire city

Downloaded from www.pdf.tube

forwarding decisions and route data between different networks.

135. (C) I, II, III, and IV
IPv4 addressing includes several specialized categories. Class A addresses use an eight-bit prefix, while Class D is dedicated to multicasting. The address 127.0.0.1 is globally recognized as the loopback address for testing local systems. To address the shortage of available addresses, Classless Inter-Domain Routing was introduced. These concepts are all fundamental to how the internet has historically managed and allocated addresses.

136. (B) Flow control
The data link layer is responsible for several critical functions to ensure reliable communication over a physical link. One of its primary roles is flow control, which manages the rate of data transmission between two adjacent nodes. By regulating the amount of data sent before an acknowledgment is received, the data link layer prevents a faster sender from overwhelming a slower receiver.

137. (D) A dedicated path is established for the duration of the connection.
A defining characteristic of a circuit-switched network is the establishment of a dedicated physical communication path between two nodes for the entire duration of their session. Once the connection is set up, all data follows this same path, and the network resources are reserved exclusively for those users. This ensures a guaranteed bit rate and constant delay, making it ideal for voice calls.

138. (A) 7
The Open Systems Interconnection model is a conceptual framework developed by ISO to standardize the functions of a telecommunication or computing system. It consists of exactly seven distinct layers: Physical, Data Link, Network, Transport, Session, Presentation, and Application. Each layer serves a specific purpose, allowing different vendors and technologies to work together by following a common set of rules and protocols.

139. (C) Router
A router is a sophisticated network device responsible for connecting multiple independent networks and forwarding data packets between them. By analyzing the destination IP address of each incoming packet and consulting its internal routing table, the router determines the most efficient path for the data to reach its final destination. Routers are essential for directing traffic across the complex global internet infrastructure.

140. (A) I, III, and V
Multiple access methods allow multiple stations to share the same physical medium. Techniques like CSMA/CA and FDMA are designed for this purpose. However, HDLC is a data link layer protocol for framing, and SMTP is an application layer protocol for email. A Virtual Private Network is a security technology. Therefore, these three are not methods for managing multiple access to a shared channel.

NODIA

141. (D) Network Layer
The network layer is responsible for the fragmentation and reassembly of data datagrams. When a large packet needs to travel through a network segment with a smaller Maximum Transmission Unit, the network layer at the router or source host breaks it into smaller fragments. These fragments are then reassembled by the network layer at the final destination host before being processed.

142. (B) The protocol to be used, the domain name/IP address, and the path of the resource.
A Uniform Resource Locator is a comprehensive address used to locate and access resources on the World Wide Web. It consists of three essential parts: the protocol used for communication, the domain name or IP address of the hosting server, and the specific path or location of the resource on that server. This structured format ensures that web browsers can find the data.

143. (C) Private Key
To create a secure digital signature that provides non-repudiation, the sender must encrypt a mathematical hash of the message using their own private key. Since the private key is known only to the sender, anyone can use the sender's public key to verify the signature. This process

communication, as the signal is only sent from the source to the receiver without any return transmission or feedback.

162. (A) Data Modulation
Data modulation is the process of converting digital information into analog signals so they can be transmitted over an analog communication medium. This involves varying one or more properties of a high-frequency carrier wave, such as its amplitude, frequency, or phase, to represent the binary ones and zeros. This technique is essential for transmitting computer data over media like radio or telephone lines.

163. (C) Data Link Layer
A network switch is a sophisticated device that operates at the data link layer of the OSI model. Its primary function is to receive data frames and forward them only to the specific port connected to the destination device. Unlike a hub, a switch uses physical MAC addresses to make intelligent forwarding decisions, which significantly improves network efficiency and reduces unnecessary traffic.

164. (D) To have a master station control when other stations can transmit.
The purpose of the polling multiple access method is to eliminate collisions on a shared medium by using a centralized control scheme. In this method, a master or primary station periodically asks each secondary station if it has data to send. A station is only allowed to transmit when it is specifically invited by the master station, ensuring orderly and collision-free communication.

165. (B) I, II, and III only
Network devices perform specific functions at different layers. Hubs broadcast traffic to all ports, while switches forward data only to the relevant destination port based on MAC addresses. Routers connect separate networks using IP addresses. However, bridges operate at the data link layer, not the network layer, to connect different network segments. Therefore, the statements about hubs, switches, and routers are correct.

166. (B) Time to Live (TTL)
The Time to Live field in the IPv4 header is a critical mechanism used to prevent packets from looping indefinitely within a network. Each router that processes the packet decrements the TTL value by at least one. If the TTL reaches zero before the packet reaches its destination, the router discards it and typically sends an error message back to the source.

167. (D) The source and destination are on different networks and the packet must go through a router.
Indirect network layer delivery occurs when the source and destination hosts are located on different logical networks. In this scenario, the sender cannot deliver the packet directly to the final host. Instead, it must send the datagram to an intermediate router, which then uses its routing capabilities to forward the packet towards the destination network, potentially across multiple intermediate network hops.

168. (A) HTTP
The Hypertext Transfer Protocol is the foundation of data communication for the World Wide Web. Its primary responsibility is to facilitate the transfer of web documents, such as HTML pages, images, and other resources, from a web server to a client's web browser. HTTP follows a request-response model, where the browser requests a resource and the server provides the corresponding data.

169. (C) It is a self-organizing, infrastructure-less network of mobile nodes.
The main characteristic of a Mobile Adhoc Network is that it is a decentralized and self-organizing network of mobile devices. Unlike traditional networks, MANETs do not require any fixed infrastructure, such as base stations or access points. Instead, each device in the network acts as a router, forwarding data for other nodes to enable communication across a dynamic and changing topology.

170. (B) I, II, III, IV, and V
Network security utilizes various mechanisms to protect data and systems. Encryption through public-key and secret-key algorithms ensures confidentiality. Digital signatures provide authentication and integrity. Firewalls filter traffic to block threats, and steganography

Click Here to Buy book on Amazon

224. (B) The efficient and dynamic allocation of computing resources (CPU, memory, storage) to users as needed.
Resource management in cloud computing is the process of efficiently and dynamically allocating physical and virtual resources to meet user demands. This involves monitoring the usage of CPU, memory, and storage across the cloud infrastructure and redistributing these resources in real-time. Effective resource management is essential for ensuring high performance, maximizing hardware utilization, and providing the elasticity that is a hallmark of cloud services.

225. (A) P-3, Q-1, R-4, S-2
Different types of malware have unique behaviors. A virus must attach itself to a host file to spread. A worm is a standalone program that can replicate itself across networks automatically. A Trojan horse disguises its malicious intent by appearing as legitimate software. Spyware is designed to secretly monitor user activity and gather personal information. Understanding these distinctions is crucial for identifying and defending against security threats.

226. (B) It allows for more efficient and flexible allocation of IP addresses.
The primary advantage of classless addressing, implemented through CIDR, is that it allows for much more efficient and flexible allocation of the IPv4 address space. Instead of being restricted to fixed classes, organizations can be assigned blocks of addresses that precisely match their actual host requirements. This significantly reduced the waste of addresses and helped slow the exhaustion of the available IPv4 address pool.

227. (D) SYN
In the TCP header, the SYN flag is used to initiate a connection between two hosts. During the first step of the three-way handshake, the client sends a segment with the SYN bit set to request synchronization of sequence numbers. This signal indicates that the client wants to establish a reliable connection and provides its initial sequence number to begin the data exchange process.

228. (A) Both A and R are true and R is the correct explanation of A.
The data link layer is responsible for the reliable delivery of frames between two adjacent nodes on a network. It achieves this node-to-node delivery by using physical MAC addresses to identify the source and destination hardware on the local link. Because this layer specifically handles the transmission across a single physical segment, the reason correctly explains the assertion regarding its core responsibility.

NODIA

229. (C) Class F
In the original classful IPv4 addressing scheme, the address space was divided into five distinct classes: A, B, C, D, and E. Classes A, B, and C were used for unicast addressing, while Class D was for multicasting and Class E for experimental or future use. Class F is not a valid or recognized class within this historical addressing system used for networking.

230. (A) Gateway
A gateway is a sophisticated networking device that connects two dissimilar networks and enables communication between them by performing protocol translation. Unlike routers, which mainly forward packets between similar networks, a gateway can translate between different communication architectures, such as connecting a local area network to a proprietary mainframe system or bridging a mobile network with the public internet using distinct protocols.

231. (D) A technology that allows transmission of data, voice, and video via a computer or any other wireless-enabled device without being connected to a fixed physical link.
Mobile computing is a broad term that refers to the technologies and infrastructure that allow people to access data and communicate while they are on the move. It involves the use of portable devices, such as smartphones and laptops, and wireless communication networks to transmit voice, video, and data. This technology enables users to remain connected to information systems regardless of their physical location.

232. (B) GSM uses TDMA and FDD, while CDMA uses spread spectrum technology.
A key technical difference between GSM and CDMA lies in their multiple access methods. GSM typically uses Time Division Multiple Access combined with Frequency Division

purpose is to gather information about the user's activities, such as web browsing habits, keystrokes, and personal data, and transmit it to a third party. Spyware can lead to significant privacy breaches and can often degrade system performance.

180. (A) Both A and R are true and R is the correct explanation of A.
General Packet Radio Service was a major advancement for GSM mobile networks. It transitioned the technology from circuit-switching to packet-switching for data. This allowed for more efficient use of the radio spectrum and enabled always-on internet connectivity for mobile devices. Because the reasoning directly supports the claim that GPRS was an important evolution, both statements are true and the explanation is correct.

181. (B) I, II, and III only
Wireless LANs utilize radio waves to provide network connectivity within homes and offices and are governed by the IEEE 802.11 standards. While they offer high flexibility, they are not inherently more secure than wired LANs. In fact, wireless signals can be intercepted more easily, making the use of strong encryption and security protocols essential for protecting data transmitted over a WLAN.

182. (A) Telephony (voice calls)
A core service provided by the Global System for Mobile Communications is telephony, which enables standard voice calls between users. GSM was originally designed as a digital system to replace older analog cellular networks, focusing on high-quality voice communication. Along with voice calls, GSM also supports other bearer and supplementary services, but telephony remains its most fundamental and widely used application.

183. (D) A vast network of uniquely identifiable, interconnected devices.
A key characteristic of the Internet of Things is the creation of a massive and vast network of uniquely identifiable, interconnected devices. These objects use sensors and software to collect and exchange data across the internet without human intervention. The goal of IoT is to integrate the physical world into computer-based systems, leading to improved efficiency, accuracy, and automation in various domains.

184. (B) It can corrupt the signal, leading to errors in the received data.
In data communication, noise refers to unwanted electrical or electromagnetic energy that interferes with the original signal as it travels through a transmission medium. The primary effect of noise is signal corruption, which can lead to errors in the received data. To maintain data integrity, communication systems use various techniques, such as shielding, filtering, and error detection and correction codes.

185. (D) I, II, III, and IV
Each network topology has unique strengths and weaknesses. Mesh networks offer high redundancy, while star topologies depend heavily on a central hub. Ring networks typically use unidirectional data flow. Bus topologies are easy to expand by connecting new devices to a single backbone cable. Understanding these characteristics is vital for designing efficient and resilient networks that meet specific connectivity and reliability requirements in environments.

186. (A) MAC Address
A Media Access Control address is a physical address represented as 48 bits, or 6 bytes. It is unique to each network interface card and is often written in a human-readable hexadecimal format, such as 00:1A:2B:3C:4D:5E. MAC addresses are used at the data link layer to ensure data is delivered to the correct hardware device on a local network segment.

187. (D) Flow Control and Error Control
The sliding window protocol is an advanced mechanism that effectively combines the functions of flow control and error control. It regulates the rate of data transmission between a sender and a receiver to prevent overwhelming the receiver's buffer. Simultaneously, it uses acknowledgments and retransmissions to ensure that any lost or corrupted data packets are identified and corrected, providing reliable communication.

188. (A) CSMA/CD
Controlled access methods are techniques used to manage how devices share a common

Downloaded from www.pdf.tube

can be used to hide sensitive information. All these techniques are integral components of a comprehensive security strategy designed to safeguard communications and data across various types of computer networks.

171. (D) PAN (Personal Area Network)
A Personal Area Network is a computer network used for data transmission among devices such as computers, telephones, and personal digital assistants within an individual's immediate workspace. PANs are typically restricted to a very small range, often just a few meters, and are commonly implemented using wireless technologies like Bluetooth to connect peripheral devices such as wireless mice, keyboards, and headsets.

172. (B) The interconnection between nodes is implemented without the use of wires.
A wireless network is a broad category of computer networks where the communication between nodes is achieved without the use of physical cables or wires. Instead, information is transmitted through the air using electromagnetic waves, such as radio waves, microwaves, or infrared signals. This technology provides flexibility and mobility, allowing devices to stay connected while moving within the coverage area of the network.

173. (C) The OSI model is a theoretical reference model, while TCP/IP is a practical, implemented protocol suite.
The relationship between the OSI and TCP/IP models is that the OSI model serves as a theoretical and standardized reference framework for understanding networking layers, whereas TCP/IP is the actual, practical suite of protocols used to implement the internet. While they share similar concepts, the TCP/IP model was developed earlier and is more streamlined, reflecting the practical requirements of real-world communications.

174. (D) Framing
The transport layer is responsible for several key functions, including process-to-process delivery, congestion control, and error control to ensure reliable end-to-end communication. However, framing is not a function of the transport layer. Framing, which involves organizing raw bits into discrete data units called frames for transmission over a physical link, is a primary responsibility of the lower data link layer.

175. (A) P-2, Q-4, R-1, S-3
Network protocols are organized into different layers. FTP is an application layer protocol for file transfers, while TCP resides at the transport layer to ensure reliability. IP is the core network layer protocol for routing and addressing. Ethernet is a standard technology that operates at the data link layer. Matching these protocols to their respective layers is essential for understanding how communication is structured.

176. (B) A network with address 10.0.0.0/8.
Classless addressing, or CIDR, allows for the flexible allocation of IP addresses using variable-length subnet masks. An address like 10.0.0.0/8 is an example of classless notation, where the /8 indicates that the first 8 bits represent the network portion. This system replaced the older classful scheme, allowing organizations to be assigned address blocks that precisely match their actual host requirements.

177. (D) Send another 1000 bytes.
In the sliding window mechanism used by TCP, the sender can transmit data up to the limit of the current window size. If the window size is 4000 bytes and the receiver acknowledges 1000 bytes, the window slides forward, allowing the sender to transmit an additional 1000 bytes. This ensures a continuous flow of data while keeping the amount of unacknowledged data.

178. (A) The process of mapping a domain name to an IP address.
In the context of the Domain Name System, resolution refers to the multi-step process used to translate a human-readable domain name into its corresponding numeric IP address. When a user enters a URL, the DNS resolution process queries various name servers to find the correct address required for the network to route the request to the appropriate destination server across the internet.

179. (C) Spyware (a type of malware)
Spyware is a specific type of malicious software that is secretly installed on a user's computer without their knowledge or consent. Its primary

calculated for each individual row and each individual column. By sending these additional parity bits along with the data, the receiver can identify both the row and the column where an error occurred, significantly improving the detection capability.

198. (C) VLANs allow a network administrator to logically segment a network, improving security and managing broadcast traffic.
A Virtual LAN is a technology that enables logical network segmentation. It allows administrators to group devices into separate broadcast domains regardless of their physical location on the switches. This approach improves security by isolating sensitive departments and enhances network performance by limiting the scope of broadcast traffic. VLANs are highly flexible and can be managed through software configurations rather than physical wiring.

199. (D) To verify the integrity of the packet's header.
In an IPv4 datagram, the checksum field is used exclusively to verify the integrity of the header information. As the packet travels through routers, the checksum is checked at each hop to ensure that no errors were introduced in critical fields like the destination IP address. Because routers often modify the header, the checksum is recalculated and updated by each router.

200. (C) P-3, Q-4, R-1, S-2
Mobile communication terms describe different technologies and services. GSM is a global standard using TDMA, while GPRS provides packet-switched data on these networks. MANETs are infrastructure-less networks formed by mobile nodes. SMS is a widely used service for sending short text messages. Correctly identifying these terms and their descriptions is key to understanding the architecture and functionality of modern mobile systems.

201. (B) IaaS
Infrastructure as a Service is the most appropriate cloud model for organizations that need to rent fundamental computing resources like virtual machines and storage. In this model, the provider manages the underlying hardware, but the user has full control over the operating system and software. This flexibility is essential for hosting legacy applications that require specific configurations and direct management of the server.

202. (A) To function as a repeater in the sky.
The primary function of a communication satellite is to act as a sophisticated repeater located in space. It receives signals from an earth station, amplifies or regenerates them, and then retransmits them back to another earth station or multiple receivers. This allows for communication over extremely long distances and bypasses physical barriers that would block terrestrial-based line-of-sight transmission methods or cables.

203. (C) Short Message Service (SMS)
Short Message Service is a key mobile communication service offered under the GSM standard. It allows for the exchange of short text messages between mobile devices. Originally developed as part of the digital cellular standards, SMS has become a ubiquitous tool for person-to-person communication and automated alerts due to its simplicity, reliability, and support across virtually all mobile networks and device types.

204. (D) Storing data on a network of data servers accessible through the internet.
Cloud storage is a model of data storage where digital information is stored on a distributed network of servers managed by a cloud provider. Users access and manage their data over the internet rather than on their local computer's hard drive. This approach provides several benefits, including high data availability, easy accessibility from multiple devices, and improved scalability for both individuals and organizations.

205. (C) The first 24 bits represent the network portion of the address.
In CIDR notation, the number following the slash represents the length of the network prefix. For the address 172.16.10.5/24, the /24 signifies that the first 24 bits are used to identify the network, while the remaining 8 bits are used to identify individual hosts within that network. This allows for more precise and efficient division of the IP address space compared to classes.

206. (A) Logical addressing and routing
The network layer, which is the third layer in the OSI model, is primarily responsible for logical

Click Here to Buy book on Amazon

addressing and routing. Its goal is to ensure that data packets are delivered from the source host to the destination host across interconnected networks. To achieve this, the network layer uses logical IP addresses and determines the best path for each packet using various routing algorithms.

207. (D) DNS - Maps physical to logical address
DNS is incorrectly described as mapping physical addresses to logical addresses. In reality, the primary function of the Domain Name System is to resolve human-readable domain names into the numeric logical IP addresses used by computers. Mapping logical addresses to physical MAC addresses is handled by the Address Resolution Protocol, while mapping physical to logical addresses is generally not a primary function of DNS.

208. (B) Public-key uses a pair of keys (public and private), while secret-key uses a single shared key.
The fundamental difference between these two cryptographic methods lies in their key management. Secret-key, or symmetric, cryptography relies on a single shared key for both encryption and decryption. Public-key, or asymmetric, cryptography uses a pair of mathematically related keys: a public key that can be shared with anyone for encryption, and a private key kept secret by the owner for decryption.

209. (C) Bearer services, Teleservices, and Supplementary services.
The service architecture of the Global System for Mobile Communications is organized into three main categories. Bearer services provide the capacity to transmit signals between access points. Teleservices provide the full capacity for communication between users, such as voice calls and SMS. Supplementary services offer additional features like call forwarding and caller ID, enhancing the basic communication experience for mobile users.

210. (A) Both A and R are true and R is the correct explanation of A.
The Internet of Things significantly impacts network traffic because it involves a massive number of interconnected devices. These devices, ranging from smart appliances to industrial sensors, often collect and transmit data continuously. Even if individual messages are small, the sheer volume of devices creates a massive amount of total data traffic. This logically explains why the expansion of IoT is a major consideration for net workers.

211. (D) I, II, III, and IV
TCP and UDP are the two primary transport layer protocols with different characteristics. TCP provides reliable, ordered delivery and congestion control, making it slower but safer. UDP is faster due to its lower overhead, making it ideal for real-time applications like video streaming where speed is prioritized over perfect reliability. Both protocols are essential for different types of network communication based on the application's needs.

212. (B) Simplex
A standard keyboard and monitor connection typically uses simplex communication. In this mode, data travels in only one direction: from the keyboard to the computer and from the computer to the monitor. There is no return path for data to travel back to the source device over the same channel, as the communication is designed to be purely unidirectional between the two components.

213. (C) Source MAC address
A network switch builds and maintains its forwarding table by examining the source MAC address of every incoming data frame. When a frame arrives on a specific port, the switch records the MAC address and the port it came from. This allows the switch to learn which devices are connected to which ports, enabling it to forward future frames only to the correct destination.

214. (D) Presentation Layer
The presentation layer, which is the sixth layer of the OSI model, is specifically concerned with the syntax and semantics of the information exchanged between two communicating systems. Its role is to ensure that data sent from the application layer of one system is readable by the application layer of another, handling tasks like data formatting, encryption for security, and compression for efficiency.

215. (B) I, III, and IV only

NODIA

IPv6 was developed to replace IPv4, offering several improvements. It features a massive 128-bit address space and supports address autoconfiguration. Unlike IPv4, fragmentation in IPv6 is handled only by the source host to reduce the processing burden on routers. Additionally, the IPv6 header does not include a checksum, as error checking is handled at other layers to further improve routing efficiency in the network.

216. (B) Spans a large geographical area, often connecting multiple LANs.
A Wide Area Network is a computer network that extends over a large geographical distance, such as across cities, countries, or continents. WANs are typically used to interconnect several local area networks so that users in one location can communicate with users in another. They often utilize various transmission media, including leased lines, fiber-optic cables, and satellite links, to span these vast distances.

217. (A) Foreign Agent
In the Mobile IP protocol, when a mobile user moves away from their home network, they connect to a new network that uses a Foreign Agent. The Home Agent in the user's home network works with this Foreign Agent to forward data to the mobile device's current location. This system allows the device to stay reachable even as it changes its point of attachment.

218. (C) Defines the performance, availability, and responsibilities guaranteed by the provider.
A Service Level Agreement is a critical document in cloud computing as it formalizes the relationship between the service provider and the customer. It explicitly defines the level of service to be provided, including guarantees on system uptime, performance metrics, and technical support. The SLA also outlines the responsibilities of both parties and the remedies or penalties if the agreed-upon standards are not met.

219. (D) Digitization (e.g., Pulse Code Modulation)
Digitization is the process of converting analog information, such as voice or video, into a digital format consisting of binary zeros and ones. One common method for achieving this is Pulse Code Modulation, which involves sampling the analog signal at regular intervals and quantizing the measured values into discrete digital levels. This allows the information to be processed, stored, and transmitted by computer systems.

220. (A) I, II, and III only
Cloud service models represent the different ways that cloud computing resources are provided to customers. The three primary models are Software as a Service, Platform as a Service, and Infrastructure as a Service. Public Cloud and Private Cloud, on the other hand, are cloud deployment models, which describe how the cloud infrastructure is owned, managed, and made available to users in various environments.

221. (B) Guard bands
In Frequency Division Multiplexing, multiple signals share a single communication channel by being assigned different frequency bands. To prevent interference between these adjacent channels, unused frequency strips called guard bands are placed between them. These buffers ensure that the signals do not overlap or bleed into each other, maintaining the clarity and integrity of each individual communication stream across the shared medium.

NODIA

222. (A) IMAP
The Internet Message Access Protocol is the standard protocol used for retrieving and managing email stored on a remote server. Unlike POP3, which often downloads and removes mail, IMAP keeps the messages on the server and allows users to create, delete, and organize folders. This ensures that the user's email environment remains synchronized across all their devices, such as phones, tablets, and computers.

223. (D) Encryption
A Virtual Private Network creates a secure and private tunnel across a public network infrastructure by using robust encryption. By encrypting the data before it is sent and decrypting it at the receiving end, the VPN ensures that even if the information is intercepted, it remains unreadable to unauthorized parties. This provides confidentiality and integrity for sensitive communications over less secure networks.

medium by preventing collisions. Reservation, polling, and token passing are all examples of controlled access, as they determine when a station is allowed to transmit. However, CSMA/CD is a contention-based method where stations compete for the channel and detect collisions when they occur, rather than being managed.

189. (C) The network portion
Routers use the network portion of an IPv4 address to make forwarding decisions. When a router receives a packet, it ignores the specific host part and looks at the network bits to determine which of its outgoing interfaces will lead closer to the destination network. This hierarchical addressing structure allows routers to efficiently direct traffic across the global internet without knowing every individual.

190. (B) I, IV, and V
The architecture of electronic mail relies on specific protocols for different tasks. The Simple Mail Transfer Protocol is used for sending messages. For retrieving emails from a server, the Post Office Protocol or the Internet Message Access Protocol is utilized. FTP and TELNET serve different purposes, such as file transfer and remote login, and are not considered primary components of the email architecture.

191. (D) Overcome geographical barriers and provide wide area coverage.
The primary purpose of communication satellites is to provide wide area coverage and overcome geographical barriers that make terrestrial communication difficult or impossible. By functioning as repeaters in space, satellites can relay signals over vast distances, connecting remote regions, ocean-going vessels, and aircraft. This makes them essential for global broadcasting, long-distance telephony, and providing internet access to underserved areas worldwide.

192. (B) Mobile IP
Mobile IP is a communication protocol designed to support mobility by allowing a user to maintain their original home IP address while moving between different networks. This ensures that active data sessions and connections are not disrupted as the device changes its point of attachment. It achieves this by using home and foreign agents to tunnel data to the device's location.

193. (C) Private Cloud, SaaS
When a company develops its own application and hosts it on its own cloud infrastructure for its employees to use via a browser, it is using a private cloud deployment model. Since the service provides a fully functional application, it is also an example of Software as a Service. This combination allows the organization to maintain high control and security over its data.

194. (D) PaaS
Platform as a Service is a cloud model where a provider offers a complete development and deployment environment, including virtual servers, operating systems, and databases. This allows developers to focus on building and managing their applications without the complexity of configuring and maintaining the underlying infrastructure. PaaS significantly speeds up the software development lifecycle and reduces the need for specialized hardware.

195. (A) 4, 1, 2, 3
Domain name resolution follows a logical hierarchy. A client first checks its local cache for the IP address. If not found, the resolver queries a root DNS server to start the process. This leads to the authoritative server for the specific domain. Finally, the client receives the correct IP address. This sequence ensures that human-readable names are efficiently translated into machine-usable addresses.

196. (B) Simplex is one-way; half-duplex is two-way but not at the same time.
The fundamental difference between simplex and half-duplex modes lies in the directionality and timing of communication. Simplex is strictly a one-way transmission, like a radio broadcast. Half-duplex allows for two-way communication, but only one party can transmit at a given time, requiring them to take turns. Full-duplex, in contrast, allows both parties to send and receive information simultaneously over the channel.

197. (A) Two-Dimensional Parity Check
Two-Dimensional Parity Check is an error detection method that organizes data bits into a table of rows and columns. A parity bit is

Downloaded from www.pdf.tube

Duplexing, where users are separated by time slots. In contrast, CDMA uses spread spectrum technology where each user is assigned a unique code, allowing multiple users to share the same frequency band simultaneously without interfering with each other's signals.

233. (C) Virtualization
The concept of a virtual server is a direct application of virtualization technology. By using a hypervisor to abstract the physical hardware, a single physical server can be divided into multiple independent virtual instances. Each virtual server operates with its own operating system and resources, allowing for more efficient use of hardware, improved isolation between users, and greater flexibility in managing server environments.

234. (D) To collect data from the physical environment.
In the context of the Internet of Things, the primary role of sensors is to act as the interface between the physical world and the digital network. Sensors are used to collect various types of real-time data from the physical environment, such as temperature, humidity, light, motion, or pressure. This information is then transmitted across the IoT network to be processed and used for decision-making.

235. (B) 3, 2, 1
Networks are often categorized by the geographical area they cover. Local Area Networks serve small spaces like homes or offices. Metropolitan Area Networks connect users within a city-sized region. Wide Area Networks span the largest distances, often connecting multiple cities, countries, or even continents. Arranging these in order of their typical physical size helps in understanding the scale and scope of network infrastructure.

236. (B) Resolution
In the Domain Name System, resolution is the process through which a hostname is mapped to its corresponding IP address. When a user requests a domain, the resolver software queries various DNS servers in a hierarchical fashion until it finds the server that has the correct address records. This essential service allows users to navigate the internet using memorable names instead of numeric IP addresses.

237. (B) IaaS
Infrastructure as a Service is the most foundational cloud service model, providing users with basic building blocks such as virtual servers, network connections, and storage. In this model, the provider manages the underlying physical infrastructure, but the customer is responsible for setting up and managing the operating system, middleware, and applications, offering the highest level of control over the computing environment.

238. (A) It gives a station permission to transmit data.
In a token ring network, the token is a special control frame that circulates continuously around the ring. Its purpose is to manage access to the shared medium and prevent collisions. A station that has data to send must wait for the token to arrive and capture it. While holding the token, the station has exclusive permission to transmit its data onto the ring.

NODIA

239. (C) Flow control prevents a fast sender from overwhelming a slow receiver; congestion control prevents the network itself from being overloaded.
While both mechanisms regulate data transmission in TCP, they serve different purposes. Flow control is an end-to-end process that prevents a sender from transmitting data faster than the receiver can process it, protecting the receiver's buffer. Congestion control is a network-wide process that monitors for signs of network overload and reduces the transmission rate to prevent the entire communication infrastructure from becoming overwhelmed.

240. (A) Both A and R are true and R is the correct explanation of A.
The User Datagram Protocol is described as unreliable because it does not provide many of the features found in TCP. It lacks mechanisms for acknowledging received data, correcting errors, or retransmitting packets that are lost during transmission. Because these missing features are the definition of an unreliable service in networking, the reasoning provided is a correct and direct explanation for the assertion.

241. (B) I, II, and III only
Firewalls are essential security tools that can be implemented as hardware, software, or a

combination of both. They act as a barrier between trusted and untrusted networks, and packet-filtering firewalls use IP addresses and port numbers to make decisions. However, firewalls cannot prevent all types of malware, as some malicious software can enter through legitimate traffic or external storage devices.

242. (A) Session Layer
The session layer, which is the fifth layer of the OSI model, is responsible for establishing, managing, and terminating dialogues between two communicating computers. It controls the synchronization of the data exchange and manages the session's duration. By providing these services, the session layer ensures that communication sessions are handled in an orderly and reliable fashion, even if there are interruptions in the network connection.

243. (C) Mobile communications
The cellular topology is a foundational concept in mobile communications. It involves dividing a large geographic area into many smaller, hexagonal or circular regions called cells. Each cell is served by its own base station and uses a specific set of radio frequencies. This design allows for frequency reuse and efficient use of the limited radio spectrum, enabling high-capacity and widespread mobile network coverage.

244. (D) The application requires guaranteed, in-order delivery of data.
UDP is often chosen over TCP for applications that require low latency or can tolerate some data loss, such as live streaming or gaming. It is also more efficient for multicasting. However, if an application requires guaranteed and in-order delivery of data, TCP is the appropriate choice. UDP does not provide these reliability mechanisms, so it would not be a valid reason to use it.

245. (D) I, II, III, and IV
The World Wide Web relies on several core technologies working together. URLs identify specific resources, and DNS translates domain names into IP addresses. HTTP is the protocol used for transferring data between servers and browsers. HTML is the standard language for creating the structure of web pages. All these components are essential for the functioning of the web as we use it today in digital communication.

246. (B) Port addresses
At the transport layer, port addresses are used to identify the specific source and destination application processes running on host computers. While IP addresses identify the physical devices, port numbers allow the transport layer to deliver data to the correct software entity, such as a web server or email client. This addressing level is essential for supporting multiple simultaneous applications on a single machine.

247. (A) An emulated server instance running on a physical host.
A virtual server is a software-based emulation of a physical computer that runs on a host server using virtualization technology. Each virtual server has its own dedicated operating system and applications, functioning independently even though it shares the underlying hardware with other virtual instances. This approach allows for more efficient resource utilization and provides flexible, on-demand computing power in cloud and data center environments.

248. (C) Confidentiality
Digital signatures are used in network security to provide authentication, integrity, and non-repudiation. They prove the sender's identity and ensure the message has not been altered. However, digital signatures alone do not provide confidentiality. Confidentiality, which ensures that only the intended recipient can read the message, must be achieved through encryption techniques, often used alongside digital signatures for a complete security solution.

249. (D) Ethernet LAN
Baseband transmission is a technique where the entire bandwidth of the transmission medium is used to carry a single digital signal at a time. A classic example of this is an Ethernet LAN, which uses baseband signaling to transmit data over twisted-pair or coaxial cables. In contrast, systems like cable television or radio use broadband transmission, which carries multiple signals simultaneously using different frequencies.

250. (D) I, II, III, IV, and V

Click Here to Buy book on Amazon

Both the OSI and TCP/IP models provide frameworks for understanding network communication. The OSI model is a theoretical seven-layer reference, while the TCP/IP suite is a practical model used for the internet. Both were developed to standardize international communications and share common layers for routing. In modern networking, the TCP/IP model is more widely implemented, making all the provided statements about them accurate.

251. (B) Full-duplex
Full-duplex communication allows data to flow in both directions at the same time. This mode offers maximum flexibility compared to simplex or half-duplex methods. It enables devices to send and receive information concurrently without waiting for the channel to become clear. This bidirectional capability is essential for efficient and modern high-speed data exchange systems across various networks.

252. (A) Bandwidth represents the theoretical capacity of a channel, while throughput is the actual successful data transfer rate.
Bandwidth represents the maximum theoretical capacity or the widest range of frequencies a transmission channel can support. In contrast, throughput measures the actual rate at which data is successfully delivered over that channel in real-world conditions. While bandwidth sets the potential limit, throughput accounts for overhead, errors, and delays that often reduce the effective transmission speed significantly.

253. (C) Simplex mode
Television broadcasting serves as a classic example of simplex communication because signals move in only one direction. The station acts as the transmitter, while the viewers' televisions act solely as receivers. Since the audience cannot send data back through the same broadcast channel, the transmission remains strictly unidirectional, functioning similarly to a one-way street for information flow.

254. (D) A sequence of discrete voltage pulses.
Digital signals are characterized by a sequence of discrete voltage pulses representing binary data. Unlike analog signals, which consist of continuous waves, digital representations use distinct levels to convey information. These discrete states allow for more robust transmission and easier error detection. Modern communication systems rely heavily on these voltage sequences to process and transmit data efficiently across various mediums.

255. (C) The time delay between sending a signal and it being received.
Latency refers to the specific time delay encountered from the moment a signal is transmitted until it is successfully received at its destination. This metric is a critical performance indicator for data communication systems. High latency can negatively impact real-time applications by causing noticeable lags. It encompasses various factors including propagation time, processing delays, and queuing within the network infrastructure.

256. (B) Baseband transmission
Baseband transmission utilizes the entire available bandwidth of a communication medium to send a single signal at any given time. This approach typically involves transmitting data in a digital format. Because the signal occupies the whole channel, only one transmission can occur simultaneously. This method is common in local area networks where high-speed communication over short distances is required.

NODIA

257. (D) To combine multiple data streams into a single signal for a shared medium.
Multiplexing is a technique used to combine multiple independent data streams into a single composite signal for transmission over a shared medium. This process ensures the efficient use of a transmission medium's total capacity. By allowing several signals to share the same physical path, multiplexing maximizes resource utilization. It is a fundamental concept in modern telecommunications for managing high-volume data traffic.

258. (A) Encoding converts digital data to digital signals, while modulation impresses information onto a carrier for analog transmission.
Encoding involves converting digital data into digital signals that are optimized for specific transmission mediums using various line coding schemes. Modulation differs by impressing information onto a carrier signal, making it suitable for analog transmission. While encoding

value of the root node.
(D) simplify the representation of the game tree itself.

9. Which of the following is considered an incorrect statement about knowledge representation?
(A) Logic provides a formal language for representing information.
(B) Semantic networks represent knowledge as a graph of interconnected nodes and links.
(C) Scripts are primarily used for representing declarative, static facts.
(D) Frames are data structures for representing stereotypical situations.

10. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : In game trees, the alpha value at a MAX node can only increase or stay the same.
Reason (R) : The alpha value represents the minimum score that the maximizing player is assured of at that node or any of its ancestors.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

11. What does 'Conceptual Dependency' (CD) theory aim to represent?
(A) The hierarchical relationship between different high-level concepts.
(B) The uncertainty associated with empirical or factual knowledge.
(C) The grammatical structure of complex natural language sentences.
(D) Sentence meaning using a small set of primitive action components.

12. In the context of knowledge representation, what is an ontology?
(A) A system for handling uncertain or incomplete information.
(B) A formal, explicit specification of a shared conceptualization.
(C) A specific instance of a frame structure.
(D) A collection of if-then rules for an expert system.

13. Consider the following statements about knowledge representation techniques:
(a) Semantic networks use nodes to represent concepts and arcs to represent relations.
(b) Frames use slots and fillers to store attributes of an object or concept.
(c) Rules in an expert system are typically in an 'IF-THEN' format.
(d) Scripts are best suited for representing common sequences of events.
(E) Ontologies primarily focus on procedural knowledge.
Which of the above statement(s) is/are correct?
(A) A, B, and D only
(B) A, B, C, and D only
(C) B, C, D, and E only
(D) A, C, and E only

14. Which component of an expert system is responsible for inferring new facts from existing data and rules?
(A) Knowledge Base
(B) User Interface
(C) Explanation Facility
(D) Inference Engine

15. What is the fundamental problem that Goal Stack Planning attempts to solve, which is addressed by non-linear planning?
(A) The Sussman Anomaly, where a subgoal solution undoes a previous one.
(B) The inability to represent complex goals in a stacked data structure.
(C) The difficulty in finding a valid initial state for a generated plan.
(D) The excessive memory usage compared to hierarchical planning models.

16. Which of the following is NOT a component of a classical planning system?
(A) A representation of the states of the world.
(B) A representation of agent emotions and beliefs.
(C) A representation of the goal to be achieved.
(D) A representation of the possible actions.

17. Which statement accurately distinguishes between linear and non-linear planning?
(A) Linear planning executes actions in parallel, while non-linear planning executes them sequentially.

Downloaded from www.pdf.tube

267. (B) Combines flow control and error control, allowing multiple frames to be sent before acknowledgment.
The Sliding Window Protocol enhances network efficiency by combining flow and error control mechanisms. It allows a sender to transmit multiple data frames before requiring an acknowledgment from the receiver. This approach is a significant improvement over simpler stop-and-wait protocols where the sender must pause after every frame. It maximizes the use of available bandwidth and reduces time spent waiting for confirmations.

268. (A) High-Level Data Link Control (HDLC)
High-Level Data Link Control is a standardized, bit-oriented protocol that provides a comprehensive solution for managing data links. It integrates several critical functions, including framing, flow control, and error control. HDLC is designed to support both point-to-point and multipoint communication, ensuring reliable data delivery across a variety of network topologies. Its structured approach makes it a fundamental protocol at the Data Link Layer.

269. (C) Flow control
Flow control is the essential function used to manage the rate of data transmission between a fast sender and a slow receiver. This mechanism prevents the receiver's buffers from being overwhelmed by an excessive influx of data. By coordinating the speeds of the two devices, flow control ensures that data is processed efficiently and prevents the loss of information due to buffer overflow.

270. (D) Detection mechanisms identify corruption, which then triggers correction or retransmission.
Error control involves a logical sequence where detection mechanisms first identify if data corruption occurred during transmission. Once an error is detected through techniques like checksums or cyclic redundancy checks, the system initiates correction or retransmission procedures. This two-step process ensures data integrity by first pinpointing problems and then applying the necessary steps to restore the original, uncorrupted information for the receiver.

NODIA

271. (C) CSMA/CD detects collisions, while CSMA/CA attempts to avoid them before they happen.
CSMA/CD and CSMA/CA differ primarily in their approach to handling collisions on a network. CSMA/CD, common in wired Ethernet, detects collisions as they happen and then retries transmission. Conversely, CSMA/CA is used in wireless environments where collision detection is difficult, so it focuses on avoiding collisions before they occur by using specific signals like acknowledgments to coordinate channel access.

272. (B) Token passing
Token passing is a controlled access method where a specific frame, called a token, circulates among the stations on a network. A station is only permitted to transmit data when it holds this token. This mechanism ensures orderly access to the shared medium and eliminates the possibility of collisions, as only one device can have the authority to send data at any time.

273. (D) CSMA/CA
While FDMA, TDMA, and CDMA are categorized as channelization protocols that divide bandwidth, time, or code space, CSMA/CA is a random access protocol. Channelization methods allow for simultaneous, collision-free transmission through division, whereas CSMA/CA focuses on avoiding collisions in environments where stations compete for immediate access. Therefore, it does not fit the definition of a protocol that pre-allocates channel resources to users.

274. (A) Polling
Polling is a controlled access protocol where a primary or central controller systematically queries each secondary station to determine if it has data to send. This method prevents collisions by ensuring that only the station specifically addressed by the controller can transmit at that moment. It provides an organized and deterministic way to manage access to a shared communication channel in various network architectures.

275. (C) Coordinate access to a shared communication channel.
The fundamental purpose of multiple access protocols is to coordinate access to a shared communication channel among several stations.

focuses on the representation of bits, modulation adapts the signal to traverse channels that may require specific frequencies or wave characteristics for delivery.

259. (C) Dividing the medium's bandwidth into separate channels.
Broadband transmission enables the simultaneous transport of multiple independent signals by dividing the available bandwidth of the medium into several separate channels. This frequency-division approach allows different data streams to travel at the same time without interfering with one another. It stands in contrast to baseband methods, which dedicate the entire medium's capacity to a single, high-speed digital transmission.

260. (A) Analog transmission.
Converting digital data into an analog signal for the purpose of sending it over a medium is classified as analog transmission. This process often involves modulation techniques where data bits are impressed upon a continuous carrier wave. This method is necessary when the physical medium is designed to carry continuous electromagnetic waves rather than discrete voltage pulses used in digital signaling.

261. (D) Port address
Port addresses are utilized at the Transport layer to identify specific application processes running on a host computer. While physical and logical addresses handle device identification and network routing, port numbers ensure that data reaches the correct software service. This addressing scheme allows multiple applications to communicate simultaneously over a single network connection by differentiating each unique data flow.

262. (B) It simplifies network design by dividing complex tasks into manageable, independent layers.
A layered architecture simplifies the complex task of networking by dividing it into smaller, more manageable components. Each layer performs a specific function and provides services to the layer above it while hiding implementation details. This modular design promotes interoperability, easier troubleshooting, and the independent development of protocols. It ensures that changes in one layer do not necessarily disrupt the entire system.

263. (C) OSI is a conceptual, seven-layer model, while TCP/IP is a more pragmatic, implementation-derived suite.
The OSI model is a conceptual seven-layer framework that serves as a theoretical standard for network protocol development. In contrast, the TCP/IP suite is a more pragmatic, implementation-derived model that serves as the foundation for the modern internet. While OSI is prescriptive, TCP/IP focuses on practical functionality and is typically described using four or five layers in most networking contexts.

264. (D) Network Layer
Logical addresses, such as IP addresses, are fundamentally associated with the Network layer in a layered network model. This layer is responsible for the host-to-host delivery of packets across potentially multiple networks. By using logical addressing, routers can determine the best path for data to travel from the source to the destination, regardless of the underlying physical hardware or MAC addresses.

265. (A) Moving data through intermediate nodes from source to destination.
Switching techniques represent the various methods used to move data through a series of intermediate nodes from a source to a final destination. These techniques facilitate communication across complex networks by determining how data is directed through the infrastructure. They are essential for ensuring that information reaches the correct recipient efficiently, whether the network uses circuit switching, packet switching, or other protocols.

266. (B) To encapsulate network layer datagrams with headers and trailers for structured transmission.
Framing is the process of encapsulating network layer datagrams into frames by adding specific headers and trailers containing control information. This function is essential for providing a structured format for the raw bitstream during transmission. By establishing these boundaries, the receiver can identify the beginning and end of each frame, facilitating synchronization and the subsequent processing of data at the link level.

Click Here to Buy book on Amazon

A key similarity between the Transmission Control Protocol and the Stream Control Transmission Protocol is that both incorporate sophisticated congestion control mechanisms. These features allow both protocols to monitor network conditions and adjust their data transmission rates accordingly to avoid saturating the network. This shared characteristic is vital for maintaining the overall health and performance of the communication infrastructure used by various applications.

285. (C) UDP, because it is fast and does not have the overhead of reliability mechanisms.
UDP is the most suitable transport protocol for real-time online games because it prioritizes low latency and speed. In gaming, timely data delivery is more important than ensuring every single packet arrives, as a slight loss of information is often unnoticeable. UDP avoids the overhead associated with the reliability and retransmission mechanisms found in TCP, providing the fast response times required for interactive play.

286. (D) To translate human-readable domain names into machine-readable IP addresses.
The primary function of the Domain Name Service is to translate human-readable domain names into machine-readable IP addresses. This hierarchical and distributed naming system allows users to access websites using easy-to-remember names instead of numerical addresses. DNS resolution is a fundamental part of internet navigation, ensuring that requests are correctly routed to the servers hosting the desired resources or services across the globe.

287. (C) SMTP
The Simple Mail Transfer Protocol is the standard protocol used by email clients to send outgoing messages to a mail server. It is also responsible for transferring email between different mail servers across the internet. While other protocols are used for retrieving or managing mail, SMTP focuses specifically on the delivery and routing aspects of electronic mail communication between senders and recipients.

288. (B) POP typically removes mail from the server after download, while IMAP manages it on the server.
A key difference between POP and IMAP lies in how they handle messages on the server. POP generally downloads messages to a single client and then deletes them from the server, making it difficult to access the same mail from multiple devices. In contrast, IMAP allows users to manage and synchronize their email directly on the server, ensuring that messages remain available across all connected clients.

289. (D) Identify resources on the World Wide Web.
A Uniform Resource Locator is defined as a specific means to identify and locate resources on the World Wide Web. It serves as a unique address for interlinked documents, images, and other digital content. By providing a standardized format for specifying resource locations, URLs enable browsers and users to navigate the vast network of information available on the web through simple links and addresses.

290. (A) IMAP
IMAP is the best protocol for users who need to access their email from multiple devices because it synchronizes actions across all of them. Since mail is managed directly on the server, reading or deleting a message on one device will be reflected on the others. This persistent connection to the server's mailbox provides a consistent experience for users who switch between phones, laptops, and desktops.

291. (C) Cryptography makes data unreadable, while steganography hides the existence of the data.
The fundamental difference between these two techniques is that cryptography focuses on making data unreadable to unauthorized parties through encryption algorithms, while steganography focuses on hiding the very existence of the data. Cryptography transforms plaintext into ciphertext to ensure confidentiality. Steganography, however, conceals information within another medium, such as an image or audio file, so that no one suspects a message is even there.

292. (B) A Firewall
A firewall acts as a security barrier between a trusted internal network and an untrusted external network, such as the internet. It monitors and controls incoming and outgoing traffic based on a predefined set of security rules. By filtering packets and blocking unauthorized access attempts, firewalls protect the integrity and safety of the private network while allowing

Downloaded from www.pdf.tube

(d) Beta represents the worst (lowest-value) choice found so far at any choice point along the path for the minimizer.

Which of the above statement(s) is/are correct?
(A) A and B only
(B) A, B, and C only
(C) A, B, C, and D
(D) B, C and D only

26. In Natural Language Processing, what is parsing?
(A) Analyzing a string of symbols to determine its grammatical structure.
(B) The process of converting text into speech.
(C) The process of determining the meaning of a sentence.
(D) The process of translating text from one language to another.

27. Which of the following deals with the meaning of words and sentences in a language?
(A) Phonology
(B) Morphology
(C) Syntax
(D) Semantic Analysis

28. The study of how context affects the interpretation of meaning in language is known as:
(A) Semantics
(B) Pragmatics
(C) Lexical analysis
(D) Syntactic analysis

29. Arrange the following typical stages of Natural Language Processing in a logical sequence:
1. Syntactic Analysis (Parsing)
2. Morphological Analysis
3. Semantic Analysis
4. Pragmatic Analysis
5. Lexical Analysis

Choose the correct sequence:
(A) 5, 2, 1, 3, 4
(B) 1, 2, 3, 4, 5
(C) 2, 5, 1, 3, 4
(D) 2, 1, 5, 4, 3

30. In the context of multi-agent systems, what is the primary difference between an agent and an object?
(A) Objects are written in C++, while agents are typically written in Java.
(B) Agents communicate with other agents, while objects interact with the whole system.
(C) Objects maintain a state and behavior, while agents only exhibit specific behaviors.
(D) Agents are proactive and autonomous, whereas objects only react to method calls.

NODIA

31. Which of the following is a key characteristic of an agent's architecture?
(A) The way the agent's percepts are mapped to its actions.
(B) The programming language used to build the agent.
(C) The physical hardware on which the agent runs.
(D) The number of lines of code in the agent's program.

32. In multi-agent systems, the aspect concerned with how agents exchange information and coordinate actions is referred to as:
(A) The agent architecture.
(B) The communicative aspect.
(C) The agent's goal state.
(D) The agent's environment type.

33. What defines a cooperative multi-agent system?
(A) Agents compete for limited resources to maximize their individual gain.
(B) Agents have conflicting goals and try to hinder each other's progress.
(C) Agents work together, sharing information and tasks to achieve a common goal.
(D) Each agent operates in complete isolation from all other agents.

34. A scenario where multiple ride-sharing cars from different companies operate in the same city, vying for passengers, is best described as a:
(A) Competitive multi-agent system.
(B) Cooperative multi-agent system.
(C) Hierarchical multi-agent system.
(D) Single-agent system.

35. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.

Assertion (A) : Logic-based knowledge

CHAPTER 10

Artificial Intelligence (AI)

1. What is the primary purpose of the Turing Test?
(A) To measure the internal computational speed and hardware efficiency of a machine.
(B) To assess if a machine's behavior is indistinguishable from that of a human.
(C) To determine the most efficient and mathematically optimal algorithm for a problem.
(D) To evaluate the raw memory storage capacity of an Artificial Intelligence system.

2. Which of the following best describes a rational agent?
(A) An agent that acts to achieve the best possible or best expected outcome.
(B) An agent that always performs the most complex action in any environment.
(C) An agent that possesses genuine human-like emotions and self-consciousness.
(D) An agent that functions solely based on a set of pre-programmed reflex actions.

3. A problem's state space representation includes all of the following components EXCEPT:
(A) The initial state of the problem.
(B) A set of possible actions or operators.
(C) The exact sequence of actions to reach the solution.
(D) The goal state or a condition for the goal.

4. Consider the following statements about heuristic search techniques:
Statement I : Heuristic search techniques use problem-specific knowledge to find a solution more efficiently.
Statement II : Heuristic search techniques are guaranteed to find the optimal solution in all cases.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

5. In the context of game playing, what does the Min-Max search algorithm aim to do?
(A) Minimize the depth of the game tree.
(B) Maximize the score for both players simultaneously.
(C) Find the optimal move assuming the opponent also plays optimally.
(D) Randomly select moves to surprise the opponent.

6. Which procedure is an optimization technique for the Min-Max algorithm?
(A) State space pruning.
(B) Alpha-beta cutoff.
(C) Goal stack planning.
(D) Breadth-first search.

7. Match the items in List I with the items in List II regarding AI approaches and concepts and select the correct option.
List I: (Concepts)
P. Turing Test
Q. Rational Agent
R. Heuristic Search
S. State Space
List II: (Descriptions)
1. Acts to maximize performance measure
2. Set of all possible configurations of a problem
3. Imitation game
4. Uses an evaluation function to guide search
(A) P-3, Q-1, R-4, S-2
(B) P-1, Q-3, R-2, S-4
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-3, S-2

8. The primary advantage of using Alpha-Beta cutoff procedures is that they:
(A) Reduce the number of game states to be evaluated in the search tree.
(B) Guarantee finding a better solution than Min-Max.
(C) Can change the final computed Min-Max

Click Here to Buy book on Amazon

legitimate communication to pass through the interface.

293. (D) The receiver's public key is used for encryption, and the receiver's private key is used for decryption.
In a public-key algorithm, sending a confidential message typically requires using the receiver's public key for encryption. Because only the receiver possesses the corresponding private key, only they can decrypt and read the message contents. This asymmetric approach eliminates the need for the sender and receiver to share a secret key in advance, making it a highly effective method for securing communications over public networks.

294. (B) Verify the authenticity and integrity of a message.
A digital signature is a mechanism enabled by public-key cryptography that is primarily used to verify the authenticity and integrity of a digital message. It allows the recipient to confirm that the message was indeed sent by the claimed sender and that it has not been altered during transmission. While it uses cryptographic principles, its main purpose is validation rather than providing full message confidentiality.

295. (C) Malicious software designed to disrupt systems or steal information.
Malware refers to various types of malicious software designed to disrupt computer operations, gain unauthorized access to systems, or steal sensitive information. These threats can manifest in many forms, including viruses, worms, and Trojans. Understanding the nature of malware is a critical part of network security, as defenders must implement tools like antivirus software and firewalls to protect systems from these harmful digital entities.

296. (D) To allow for the creation of virtual servers and the abstraction of resources.
Virtualization is a core technology in cloud computing that allows for the creation of virtual servers and the abstraction of physical hardware resources. By separating the software from the underlying physical infrastructure, virtualization enables the efficient sharing of storage, databases, and computing power among many users. This flexibility is what allows cloud providers to offer on-demand, scalable resources to their customers through various service models.

297. (A) Mobile IP
Mobile IP is a communication protocol that allows a mobile device to maintain a constant IP address while moving between different networks. This is essential for ensuring continuous connectivity and preventing service interruptions as the device transitions from one network environment to another. By managing the routing of packets to the device's current location, Mobile IP enables seamless mobility for users on various wireless infrastructures.

298. (B) SaaS (Software as a Service)
When a company subscribes to a service that provides a complete, ready-to-use software application over the internet, they are utilizing the Software as a Service model. In this cloud computing model, the provider manages the entire stack, including the application, data, and underlying infrastructure. Users access the software typically through a web browser, eliminating the need to install or maintain the application locally on their own systems.

299. (B) A network of interconnected physical devices that collect and exchange data.
The Internet of Things is described as a vast network of interconnected physical devices that collect and exchange data over the internet. These objects range from household appliances to industrial sensors, all equipped with the technology to communicate with one another and with central systems. IoT fundamentally extends the reach of digital networks into the physical world, creating new opportunities for monitoring, automation, and data analysis.

300. (A) To govern the relationship and define service terms between a provider and a user.
A Service Level Agreement is a document that governs the relationship between a cloud service provider and a user by defining the specific terms and standards of the service. It outlines the expectations for performance, availability, and reliability that the provider must meet. The SLA serves as a critical contract, ensuring that both parties understand their responsibilities and the level of support provided by the cloud service.

NODIA

(B) Linear planning assumes subgoals are independent, while non-linear planning handles subgoal interactions.
(C) Linear planning always finds the optimal plan, while non-linear planning does not.
(D) Linear planning is used for continuous spaces, while non-linear planning is for discrete spaces.

18. What is the primary data structure used in Goal Stack Planning?
(A) A stack
(B) A queue
(C) A tree
(D) A graph

19. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : Hierarchical planning decomposes a complex problem into a hierarchy of simpler subproblems.
Reason (R) : This decomposition allows the planner to focus on abstract details first and fill in specifics later, managing complexity effectively.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

20. The STRIPS planning representation consists of which three components for defining an action?
(A) Action Name, Parameters, and Effects.
(B) Initial State, Goal State, and Operator.
(C) State, Action, and Cost.
(D) Preconditions, Add List, and Delete List.

21. Identify the incorrect statement regarding Partial Order Planning (POP).
(A) POP builds a plan by adding actions and ordering constraints incrementally.
(B) POP commits to a total ordering of all actions from the very beginning.
(C) POP is a type of non-linear planning.
(D) POP addresses threats (conflicts) between actions by adding new ordering constraints.

22. Match the items in List I with the items in List II and select the correct option.
List I: (Planning Technique)
P. Goal Stack Planning
Q. Hierarchical Planning
R. STRIPS
S. Partial Order Planning
List II: (Characteristic)
1. Represents operators with Precondition, Add, and Delete lists
2. Searches through the space of plans rather than states
3. Decomposes problems into different levels of abstraction
4. Uses a LIFO structure to manage goals and subgoals
(A) P-4, Q-3, R-1, S-2
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-1, R-2, S-3
(D) P-2, Q-3, R-4, S-1

23. Which of the following is the most significant advantage of a rational agent approach over the Turing Test approach?
(A) It is easier to implement in software.
(B) It is a more general approach that can be applied to a wider range of problems.
(C) It focuses on human-like thinking, which is the ultimate goal of AI.
(D) It provides an objective and mathematically analyzable criterion for success.

24. A search algorithm is said to be 'admissible' if it is guaranteed to:
(A) return an optimal solution if one exists.
(B) find a solution in the shortest amount of time.
(C) use the least amount of memory.
(D) explore every node in the state space.

25. Consider the following statements regarding game playing search algorithms:
(a) The Min-Max algorithm performs a complete depth-first exploration of the game tree.
(b) The effectiveness of Alpha-Beta pruning is highly dependent on the order in which moves are examined.
(c) Alpha represents the best (highest-value) choice found so far at any choice point along the path for the maximizer.

Downloaded from www.pdf.tube

These protocols provide the necessary rules to avoid or manage collisions, ensuring that data can be transmitted effectively without interference. By organizing how and when devices use the medium, they maintain the efficiency and stability of broadcast networks where resources are shared by many.

276. (A) To address the inefficient allocation of IP addresses caused by fixed classes.
Classless Inter-Domain Routing was introduced to address the inefficient allocation of IP addresses caused by the fixed boundaries of the classful addressing system. Traditional classes often resulted in wasted address space when organizations were assigned more addresses than needed. CIDR allows for variable-length subnet masking, providing a more flexible and granular method for distributing the finite pool of available IPv4 addresses across the internet.

277. (D) Fragmentation
Fragmentation is the process of breaking a large Network Layer datagram into smaller pieces so that it can fit within the frame size limits of an underlying network. When a packet exceeds the maximum transmission unit of a link, it must be divided. These fragments are then sent independently and are typically reassembled at the destination to reconstruct the original, complete data packet.

278. (B) To map a logical IP address to a physical MAC address.
The Address Resolution Protocol performs the critical function of mapping a logical IP address to a physical MAC address on a local network. Since hardware devices communicate using physical addresses at the Data Link Layer, ARP is necessary to bridge the gap from the Network Layer. It allows a sender to discover the correct hardware address associated with a specific IP destination for delivery.

279. (C) The checksum field
In the IPv4 structure, the checksum field is used specifically to verify the integrity of the packet header. It helps detect if any part of the header has been corrupted during its transit through the network. While it does not protect the data payload itself, ensuring the header is error-free is vital for correct routing and processing by intermediate nodes and the final recipient.

280. (D) Vastly larger 128-bit address space.
The most significant advantage of IPv6 compared to its predecessor is its vastly larger 128-bit address space. While IPv4 is limited to approximately four billion unique addresses, the 128-bit structure of IPv6 provides an almost inexhaustible supply. This expansion is essential for supporting the growing number of internet-connected devices worldwide and for simplifying the overall address allocation and management process on the web.

281. (A) UDP
User Datagram Protocol is a connectionless transport protocol that offers a simple and unreliable service. It does not provide mechanisms for flow control, error recovery, or congestion management beyond an optional checksum. This lack of overhead makes it very fast and highly suitable for applications where speed is prioritized over perfect reliability, such as in DNS queries or live media streaming services.

NODIA

282. (B) To reduce the transmission rate in response to network traffic overload.
Congestion control in protocols like TCP and SCTP is primarily used to manage network stability by reducing the transmission rate when traffic overload is detected. When a network becomes congested, these protocols slow down data flow to prevent complete failure or excessive delays. This mechanism ensures that the network can recover from high usage periods and maintains a fair distribution of resources among users.

283. (B) Flow control
TCP utilizes a sliding window mechanism primarily to implement flow control. This ensures that a fast-sending host does not overwhelm a slower receiver by sending more data than the receiver can process or store in its buffers. By dynamically adjusting the window size, TCP manages the amount of unacknowledged data in transit, balancing the speed of transmission with the receiver's current capacity.

284. (A) Incorporate congestion control mechanisms.

representation is monotonic.

Reason (R) : In a monotonic logic system, the addition of a new axiom (fact) to the knowledge base cannot lead to the retraction of any existing conclusions.

(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

36. Which of the following statements about Fuzzy Sets is correct?
(A) An element either fully belongs to a fuzzy set or does not belong at all.
(B) Fuzzy sets are identical to classical (crisp) sets.
(C) The membership value in a fuzzy set can only be 0 or 1.
(D) Fuzzy sets allow for degrees of membership, represented by a value between 0 and 1.

37. What is the primary purpose of a fuzzy rule-based system?
(A) To perform precise numerical calculations.
(B) To model human-like reasoning in situations of uncertainty and imprecision.
(C) To guarantee optimal solutions for planning problems.
(D) To represent knowledge using strict logical formalism.

38. Consider the following statements about Fuzzy Logic:
(a) It is a multi-valued logic that allows for intermediate truth values.
(b) The core idea is to model the vagueness inherent in human language and thought.
(c) A key component of a fuzzy logic system is the 'fuzzifier', which converts crisp inputs into fuzzy sets.
(d) A 'defuzzifier' converts the fuzzy output back into a crisp value.
(E) It is less suitable for controlling complex systems than traditional binary logic.

Which of the above statement(s) is/are correct?
(A) A, B, and E only
(B) C, D, and E only
(C) A, B, C, and D only
(D) A, C, and D only

39. The process in a Genetic Algorithm where new individuals (offspring) are created from existing ones is called:
(A) Encoding
(B) Fitness Function
(C) Initialization
(D) Reproduction

40. What is the role of the 'fitness function' in a Genetic Algorithm?
(A) To evaluate how good a particular solution is in solving the problem.
(B) To generate the initial population of solutions.
(C) To combine two parent solutions to create an offspring.
(D) To introduce random changes into a solution.

41. In Genetic Algorithms, what does 'encoding' refer to?
(A) The process of selecting individuals for reproduction.
(B) The evaluation of a solution's quality.
(C) The representation of a potential solution as a chromosome (e.g., a bit string).
(D) The random modification of a chromosome.

42. Which two operators are primarily responsible for exploration and exploitation in Genetic Algorithms?
(A) Crossover and Mutation
(B) Selection and Initialization
(C) Encoding and Decoding
(D) Fitness Calculation and Elitism

43. Match the items in List I with the items in List II and select the correct option.

List I: (Genetic Algorithm Operator)
P. Selection
Q. Crossover
R. Mutation
S. Encoding

List II: (Function)
1. Combines genetic material from two parents
2. Represents a solution in a specific format
3. Introduces random variation to maintain diversity
4. Favors fitter individuals for reproduction

(A) P-1, Q-2, R-3, S-4
(B) P-4, Q-1, R-3, S-2

Click Here to Buy book on Amazon

(A) The use of network protocols for communication.
(B) The autonomous and goal-directed nature of the components (agents).
(C) The physical separation of computational nodes.
(D) The ability to handle multiple threads of execution.

81. In the context of Genetic Algorithms, what is the 'mutation' operator's main purpose?
(A) To combine the best parts of two good solutions.
(B) To ensure the fittest individuals always survive to the next generation.
(C) To maintain genetic diversity in the population and prevent premature convergence.
(D) To evaluate the quality of a given solution.

82. The statement A multilayer perceptron without a hidden layer is equivalent to a single perceptron is:
(A) Always true.
(B) Always false.
(C) True only if the output layer uses a linear activation function.
(D) True only if the output layer uses a step activation function.

83. Which of the following is an example of an unsupervised learning algorithm?
(A) Backpropagation
(B) Reinforcement Learning
(C) Self-organizing Maps (SOM)
(D) Perceptron Learning Rule

84. Identify the incorrect statement about Partial Order Planning (POP).
(A) POP begins with a minimal plan (start and finish actions).
(B) It adds ordering constraints only when necessary to resolve conflicts.
(C) It maintains a set of causal links representing dependencies between actions.
(D) It always produces a plan that is a strictly linear sequence of actions.

85. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : Fuzzy logic is useful for modeling systems with imprecise or vague data.
Reason (R) : Fuzzy logic uses degrees of membership in sets, rather than strict binary membership, allowing it to represent concepts like 'tall' or 'hot'.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

86. What is the fundamental difference between the 'state space representation' and 'knowledge representation'?
(A) State space defines the structure of a problem for search, while knowledge representation is about storing and reasoning with facts about the world.
(B) State space is for planning, while knowledge representation is for learning.
(C) State space is always a graph, while knowledge representation is always logic-based.
(D) There is no fundamental difference; they are synonymous terms.

87. In the context of game playing, when is the Min-Max algorithm guaranteed to find the optimal move?
(A) Only if the game is finite and has two players with opposing goals.
(B) Only if the search is allowed to proceed to the terminal nodes of the game.
(C) Only if an alpha-beta cutoff is used.
(D) Both (A) and (B) are required.

88. Consider the following statements about various AI topics:
(a) STRIPS is a formal language for representing planning problems.
(b) Conceptual Dependency provides primitives for representing actions and states.
(c) A Hopfield Network is a recurrent neural network.
(d) A rational agent may not always be successful, but it does the right thing based on its knowledge.

NODIA

(A) Thinking humanly.
(B) Acting humanly.
(C) Thinking rationally.
(D) Acting rationally.

55. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : The STRIPS representation for planning operators was a significant step in AI planning.
Reason (R) : By specifying actions only in terms of what they change (add/delete lists), STRIPS avoids the frame problem of having to specify everything that does not change.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

56. Which of the following is NOT a technique for handling uncertainty in knowledge for expert systems?
(A) Certainty Factors
(B) Fuzzy Logic
(C) Propositional Logic
(D) Bayesian Probability

57. A 'frame' in knowledge representation is a structure that includes slots for storing information. What is a 'filler' in this context?
(A) The name of the frame itself.
(B) A procedure attached to a slot (a demon).
(C) A link to another related frame.
(D) The specific value or data that is placed into a slot.

58. What is the main difference between Linear Planning and Non-Linear Planning?
(A) Linear planners produce a strictly ordered sequence of actions, while non-linear planners produce a partially ordered set.
(B) Linear planners are faster, but non-linear planners are more accurate.
(C) Linear planners use stacks, while non-linear planners use queues.
(D) Linear planners cannot handle goals with multiple sub-goals.

NODIA

59. Match the items in List I with the items in List II and select the correct option.
List I: (NLP Task)
P. Semantic Analysis
Q. Pragmatic Analysis
R. Syntactic Analysis
S. Morphological Analysis
List II: (Description)
1. Analyzing sentence structure according to a formal grammar
2. Identifying the role of a word in a sentence based on its form (e.g., plural, tense)
3. Deriving the meaning of a sentence, often independent of context
4. Understanding the intended meaning considering the context and situation
(A) P-3, Q-4, R-1, S-2
(B) P-4, Q-3, R-2, S-1
(C) P-3, Q-1, R-4, S-2
(D) P-2, Q-4, R-1, S-3

60. What is a defining characteristic of a 'competitive' multi-agent system?
(A) Agents share a common utility function.
(B) The success of one agent often comes at the expense of another.
(C) Communication between agents is always open and truthful.
(D) The system's overall goal is prioritized over individual agent goals.

61. Which statement about fuzzy sets and their properties is incorrect?
(A) The union of two fuzzy sets is typically defined by the maximum of their membership functions.
(B) The intersection of two fuzzy sets is typically defined by the minimum of their membership functions.
(C) The support of a fuzzy set includes elements with a membership grade of exactly 0.
(D) The complement of a fuzzy set A is defined by $1 - \mu A(x)$.

62. Consider the following statements regarding the working principle of Genetic Algorithms:
(a) They are inspired by the process of natural selection.
(b) They typically work with a population of candidate solutions, not a single solution.
(c) They use probabilistic transition rules, not

(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-4, R-1, S-3

44. Which learning paradigm in Artificial Neural Networks involves learning from labeled data?
(A) Unsupervised Learning
(B) Supervised Learning
(C) Reinforcement Learning
(D) Heuristic Learning

45. The Backpropagation algorithm is most commonly associated with which type of neural network?
(A) Perceptron
(B) Hopfield Network
(C) Multilayer Perceptron (MLP)
(D) Self-organizing Map

46. What type of problem is a single Perceptron capable of solving?
(A) Any non-linear classification problem.
(B) Any regression problem.
(C) Problems requiring memory, like time-series prediction.
(D) Problems that are linearly separable.

47. Identify the correct statement about Unsupervised Learning in Artificial Neural Networks.
(A) The goal is to discover hidden patterns or structures in unlabeled data.
(B) It requires a teacher signal to guide the learning process.
(C) Backpropagation is the primary algorithm used.
(D) The network is trained using rewards and punishments.

48. Which of the following is a characteristic of a Hopfield Network?
(A) It is a feedforward network.
(B) It acts as a content-addressable memory system.
(C) It is used for dimensionality reduction and clustering.
(D) It uses the backpropagation algorithm for training.

49. Arrange the following basic steps of a Genetic Algorithm in the correct sequence:
1. Selection
2. Initialization (create initial population)
3. Evaluation (calculate fitness)
4. Crossover and Mutation (reproduction)

Choose the correct sequence:
(A) 1, 2, 3, 4
(B) 2, 1, 4, 3
(C) 2, 3, 1, 4
(D) 3, 2, 1, 4

50. What kind of learning involves an agent interacting with an environment and receiving rewards or penalties for its actions?
(A) Supervised Learning
(B) Unsupervised Learning
(C) Reinforcement Learning
(D) Transfer Learning

51. Consider the following statements about Artificial Neural Networks:
(a) A multilayer perceptron (MLP) can solve non-linearly separable problems.
(b) Self-organizing maps (SOMs) are a type of supervised learning network.
(c) The backpropagation algorithm calculates the gradient of the loss function with respect to the network's weights.
(d) A single perceptron uses a step function as its activation function.

Which of the above statement(s) is/are correct?
(A) A and B only
(B) A, C, and D only
(C) B and D only
(D) A, B, C, and D

52. In knowledge representation, what is a 'script' primarily designed to represent?
(A) A hierarchical classification of objects.
(B) A stereotyped sequence of events in a particular context.
(C) A set of logical rules for deduction.
(D) A graph of related concepts.

53. In a semantic network, what do the arcs typically represent?
(A) Concepts or objects.
(B) Relationships between concepts.
(C) The strength of a belief.
(D) Procedural knowledge.

54. The 'Turing Test' is an example of which approach to AI?

Downloaded from www.pdf.tube

(B) The receiver of the message.
(C) The content of the message.
(D) The programming language of the sender agent.

72. In Genetic Algorithms, the 'Crossover' operator is analogous to what biological process?
(A) Asexual reproduction
(B) Mutation
(C) Sexual reproduction
(D) Natural selection

73. The 'Sussman Anomaly' is a classic problem that demonstrates the limitations of which planning technique?
(A) Hierarchical Planning
(B) Linear Planning (like early Goal Stack Planning)
(C) Partial Order Planning
(D) Non-Linear Planning

74. Identify the correct pair of concept and its description.
(A) State Space Representation : A genetic algorithm operator
(B) Conceptual Dependency : A framework for representing the meaning of sentences
(C) Alpha-Beta Cutoff : A technique for unsupervised learning
(D) Backpropagation : A method for knowledge representation

75. Consider the following statements about AI concepts:
(a) A rational agent's action depends on its percept sequence, built-in knowledge, and goals.
(b) Min-Max search is a recursive algorithm.
(c) An ontology provides a vocabulary and a set of assumptions about the meaning of terms in a domain.
(d) Goal Stack Planning is a form of hierarchical planning.
Which of the above statement(s) is/are correct?
(A) A, B, and C only
(B) B, C, and D only
(C) A, C, and D only
(D) A, B, and D only

76. What is the primary function of a 'Hopfield Network'?
(A) To classify linearly separable data.
(B) To perform dimensionality reduction.
(C) To act as an auto-associative memory, retrieving a full pattern from a partial or noisy input.
(D) To learn from a sequence of rewards and punishments.

77. Which of the following is the best description of a 'heuristic function' in the context of search algorithms?
(A) A function that calculates the exact cost from a node to the goal.
(B) A function that estimates the cost from a node to the goal.
(C) A function that determines if a node is a goal state.
(D) A function that computes the path cost from the start node to the current node.

78. The output of a parsing technique in Natural Language Processing is typically a:
(A) parse tree or syntax tree.
(B) sequence of tokens.
(C) summary of the text.
(D) translated version of the text.

79. Match the items in List I with the items in List II and select the correct option.
List I: (AI Technique)
P. Genetic Algorithm
Q. Expert System
R. Fuzzy Logic
S. NLP Parser
List II: (Problem)
1. Diagnosing a specific type of machine failure based on symptoms.
2. Controlling a home appliance like a washing machine based on vague inputs like 'dirtiness'.
3. Optimizing the layout of components on a microchip.
4. Analyzing the grammatical correctness of a sentence.
(A) P-1, Q-3, R-4, S-2
(B) P-3, Q-1, R-2, S-4
(C) P-3, Q-2, R-1, S-4
(D) P-4, Q-1, R-3, S-2

80. What distinguishes a multi-agent system from a simple distributed system?

NODIA

(B) Agents are considered autonomous.
(C) An object must be invoked by another entity to perform an action.
(D) An agent cannot encapsulate state, only behavior.

99. Match the items in List I with the items in List II and select the correct option.
List I: (Method)
P. Logic
Q. Semantic Network
R. Script
S. Frame
List II: (Basic Element)
1. Slot and Filler
2. Node and Arc
3. Scene and Event
4. Predicate and Term
(A) P-4, Q-2, R-3, S-1
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-3, R-1, S-2
(D) P-2, Q-3, R-4, S-1

100. In a fuzzy set representing 'tall people', a person with a height of 6'2 might have a membership value of 0.9, while a person with a height of 5'8 might have a value of 0.4. This illustrates which property of fuzzy sets?
(A) They are identical to probability distributions.
(B) They allow for graded or partial membership.
(C) They only allow binary membership.
(D) They cannot be used for continuous variables like height.

101. Consider the following statements regarding learning paradigms in Artificial Neural Networks:
(a) Supervised learning uses labeled data to train a model.
(b) Unsupervised learning aims to find patterns in unlabeled data.
(c) Reinforcement learning involves an agent learning through rewards and penalties.
(d) The backpropagation algorithm is an example of reinforcement learning.
Which of the above statement(s) is/are correct?
(A) A and B only
(B) A, B, and C only
(C) C and D only
(D) A, B, C, and D

102. The role of the beta value in alpha-beta pruning is to represent:
(A) the best possible score (lowest value) for the minimizing player found so far in the search.
(B) the best possible score for the maximizing player.
(C) the depth of the current search tree.
(D) the probability of winning from a given node.

103. Which knowledge representation formalism uses if-then structures and is the backbone of most expert systems?
(A) Frames
(B) Scripts
(C) Rules
(D) Semantic Networks

104. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : A single Perceptron cannot solve the XOR problem.
Reason (R) : The XOR problem is not linearly separable, and a single Perceptron can only form a linear decision boundary.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

105. In the context of planning, a 'threat' in Partial Order Planning occurs when:
(A) a planner cannot find any applicable action.
(B) an action added to the plan could potentially negate a required condition (causal link) of another action.
(C) an action's precondition is not met.
(D) the goal state is unreachable from the initial state.

106. The main purpose of a 'communicative aspect' in a multi-agent system is to facilitate:
(A) individual agent computation.
(B) coordination, cooperation, and negotiation among agents.
(C) user interaction with the system.

Click Here to Buy book on Amazon

(E) The Turing test is a definitive and final measure of true machine consciousness.

Which of the above statement(s) is/are correct?
(A) A, B, and E only
(B) C, D, and E only
(C) A, B, C, and D only
(D) A, B, D, and E only

89. Which parsing technique builds a parse tree from the top (the start symbol) down to the leaves (the words)?
(A) Top-down parsing
(B) Bottom-up parsing
(C) Chart parsing
(D) Shift-reduce parsing

90. In a cooperative multi-agent system, what is a primary challenge?
(A) Ensuring effective coordination and communication to achieve the common goal.
(B) Preventing agents from sharing information.
(C) Designing agents that maximize their individual rewards.
(D) Fostering a competitive environment to drive innovation.

91. Arrange the following Artificial Neural Network concepts from simplest to most complex in terms of structure:
1. Multilayer Perceptron (MLP)
2. Perceptron
3. Deep Neural Network (many hidden layers)

Choose the correct sequence:
(A) 2, 1, 3
(B) 1, 2, 3
(C) 2, 3, 1
(D) 3, 1, 2

92. What does the 'Add List' in a STRIPS operator specify?
(A) The conditions that must be true before the operator can be applied.
(B) The facts that become true in the world after the operator is applied.
(C) The facts that are no longer true after the operator is applied.
(D) The cost of applying the operator.

93. The concept of an 'agent architecture' in multi-agent systems primarily refers to:
(A) the physical hardware configuration of the system.
(B) the design that maps an agent's percepts to its actions.
(C) the communication protocol used between agents.
(D) the user interface of the multi-agent system.

94. Which of the following is a key characteristic of a 'competitive' multi-agent system?
(A) Converting a crisp input into a fuzzy set.
(B) Applying fuzzy rules to the fuzzy input sets.
(C) Combining the results of multiple fuzzy rules.
(D) Converting the final fuzzy output set into a single crisp value.

95. Consider the following statements:
Statement I : The fitness function in a Genetic Algorithm must accurately reflect the quality of a solution in the context of the problem.
Statement II : A poorly designed fitness function can lead the GA to converge on suboptimal or nonsensical solutions.
Which of the following is correct?
(A) Both Statement I and Statement II are true.
(B) Both Statement I and Statement II are false.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

NODIA

96. A key feature of 'Hierarchical Planning' is that it:
(A) plans at different levels of abstraction, refining high-level plans into more detailed ones.
(B) avoids the use of preconditions and effects.
(C) guarantees finding the shortest possible plan in all cases.
(D) is a form of linear planning that cannot handle subgoal interactions.

97. In Natural Language Processing, 'semantic analysis' is primarily concerned with:
(A) the structure of sentences.
(B) the meaning of sentences.
(C) the context of conversation.
(D) the individual sounds of a language.

98. Which of the following statements about agents and objects is incorrect?
(A) Objects encapsulate state and behavior.

(E) The backpropagation algorithm is used in supervised learning.

Which of the above statement(s) is/are correct?
(A) A, B, C, and D only
(B) B, C, D, and E only
(C) A, C, D, and E only
(D) A, B, D, and E only

116. An agent's behavior is determined by its `agent = architecture + program` function. What does the 'architecture' part refer to?
(A) The physical computing device.
(B) The machinery that the agent executes on, providing percepts and receiving actions.
(C) The programming language used.
(D) The agent's knowledge base.

117. In a competitive multi-agent system, agents' goals are often:
(A) identical and shared.
(B) independent and unrelated.
(C) conflicting, such that one agent's gain is another's loss.
(D) always communicated truthfully.

118. Which statement about the 'encoding' step in a Genetic Algorithm is true?
(A) It involves evaluating the fitness of each individual.
(B) It is the final step of the algorithm.
(C) It determines how a potential solution is represented (e.g., as a binary string or a real-valued vector).
(D) It is another name for the crossover operation.

119. A Perceptron is a single-layer neural network, whereas a Multilayer Perceptron (MLP) has:
(A) one or more hidden layers between the input and output layers.
(B) multiple output neurons.
(C) a recurrent connection.
(D) no activation function.

120. Self-Organizing Maps (SOMs) are an example of which learning paradigm?
(A) Supervised learning
(B) Unsupervised learning
(C) Reinforcement learning
(D) Semi-supervised learning

121. What is the key characteristic of a 'rational agent'?
(A) It selects actions to maximize its expected performance measure.
(B) It always chooses the action that is guaranteed to succeed.
(C) It possesses a perfect model of its environment.
(D) It perfectly mimics human decision-making.

122. Which of the following is NOT a heuristic search technique?
(A) Breadth-First Search
(B) A* Search
(C) Greedy Best-First Search
(D) Iterative Deepening A*

123. The function of the 'Delete List' in a STRIPS operator is to specify:
(A) the preconditions of the operator.
(B) the new facts that are added to the state.
(C) the facts from the current state that are no longer true after the operator is applied.
(D) the operator that should be deleted from the plan.

124. An 'expert system' is a computer system that emulates the decision-making ability of a human expert. What are its two main components?
(A) A user interface and a database.
(B) A knowledge base and an inference engine.
(C) A compiler and an interpreter.
(D) A CPU and RAM.

125. Consider the following statements about Knowledge Representation:
(a) A semantic network is a graphical representation of knowledge.
(b) A frame is a data structure representing a stereotyped situation.
(c) A script represents a sequence of events.
(d) Conceptual Dependency theory uses a small number of primitive actions.

Which of the above statement(s) is/are correct?
(A) A and B only
(B) A, B, and C only
(C) C and D only
(D) A, B, C, and D

126. The backpropagation algorithm is used to:
(A) find the shortest path in a graph.

Downloaded from www.pdf.tube

(D) hardware resource management.

107. The 'fitness function' in a genetic algorithm is a measure of:
(A) how well a candidate solution solves the problem.
(B) the genetic diversity of the population.
(C) the computational time taken to generate a solution.
(D) the rate of mutation.

108. Arrange the following steps of a fuzzy inference system in the correct order:
1. Fuzzification (input)
2. Defuzzification (output)
3. Rule evaluation (inference)
4. Aggregation of rule outputs

Choose the correct sequence:
(A) 1, 3, 4, 2
(B) 1, 2, 3, 4
(C) 3, 1, 4, 2
(D) 2, 1, 3, 4

109. What is a key limitation of the Min-Max algorithm that Alpha-Beta pruning addresses?
(A) Its inability to handle games with more than two players.
(B) Its requirement of a heuristic evaluation function.
(C) The exponential growth of the search space it needs to explore.
(D) Its failure to find the optimal move in deterministic games.

110. A state space representation is NOT suitable for problems where:
(A) the sequence of actions is the solution.
(B) the problem is about representing declarative knowledge with complex relationships.
(C) the problem can be broken down into a series of state transitions.
(D) the states are discrete and finite.

111. The 'pragmatics' level of Natural Language Processing analysis is required to understand the sentence, Can you open the window? as a request rather than a question about ability. This is because pragmatics deals with:
(A) context and intent.
(B) grammatical structure.
(C) literal meaning.
(D) word formation.

112. Match the items in List I with the items in List II and select the correct option.
List I: (Concept)
P. STRIPS
Q. Hopfield Network
R. Conceptual Dependency
S. Rational Agent
List II: (Purpose)
1. A model of an agent that acts to achieve the best outcome
2. A formal language for representing planning operators
3. A network model for content-addressable memory
4. A theory for representing the meaning of natural language sentences
(A) P-1, Q-2, R-3, S-4
(B) P-2, Q-3, R-4, S-1
(C) P-2, Q-4, R-1, S-3
(D) P-4, Q-3, R-2, S-1

NODIA

113. Which of the following statements about different types of planning is correct?
(A) Hierarchical planning involves refining abstract plans into concrete steps.
(B) Goal Stack Planning is a type of non-linear planning.
(C) Partial Order Planning commits to a full ordering of actions at the start.
(D) Linear planning is immune to the problem of conflicting subgoals.

114. The 'working principle' of a Genetic Algorithm is based on:
(A) gradient descent optimization.
(B) the principles of natural evolution and survival of the fittest.
(C) logical deduction and inference.
(D) dividing a problem into smaller, independent subproblems.

115. Consider the following statements related to AI:
(a) The Turing Test is a test of a machine's ability to act rationally.
(b) An ontology defines a common vocabulary for a domain.
(c) Min-Max search is a depth-first search on the game tree.
(d) A fuzzy set allows for degrees of membership.

deterministic ones.
(d) They require the problem's search space to be differentiable.
(E) Crossover and mutation are the primary operators for generating new solutions.
Which of the above statement(s) is/are correct?
(A) A, B, C, and D only
(B) A, B, C, and E only
(C) B, C, D, and E only
(D) A, B, D, and E only

63. The 'Backpropagation' algorithm in ANNs is essentially a method for:
(A) initializing the weights of the network randomly.
(B) propagating the error backwards through the network to update the weights.
(C) selecting the optimal number of hidden layers.
(D) clustering input data in an unsupervised manner.

64. A Self-Organizing Map (SOM) is a type of neural network that is primarily used for:
(A) supervised classification.
(B) reinforcement learning tasks.
(C) dimensionality reduction and visualization of high-dimensional data.
(D) storing and retrieving patterns like a Hopfield network.

65. Arrange the following steps of the STRIPS planning algorithm in order:
1. Push the goal onto the stack.
2. If the top of the stack is a compound goal, push its subgoals onto the stack.
3. If the top of the stack is an action, pop it, and push its preconditions onto the stack.
4. If the top of the stack is a single goal satisfied by the current state, pop it.

Choose the correct sequence:
(A) 1, 2, 3, 4
(B) 1, 4, 2, 3
(C) 4, 3, 2, 1
(D) 2, 4, 1, 3

66. What is the alpha-beta cutoff procedure fundamentally based on?
(A) A probabilistic estimation of the best move.
(B) The realization that a branch of the game tree can be pruned if a move is found that is worse than a previously examined move.
(C) A method to increase the depth of the search in promising areas.
(D) The reordering of the game tree to make it balanced.

67. Which knowledge representation scheme is most suitable for representing a common sequence of events, such as going to a restaurant?
(A) Semantic Network
(B) Frame
(C) Logic
(D) Script

68. The Rational Agent approach defines intelligence in terms of:
(A) Maximizing a performance measure given the available information.
(B) Mimicking human thought processes.
(C) Passing a test of linguistic competence.
(D) The ability to solve any computable function.

69. In a fuzzy rule-based system, the process of converting a crisp input value (e.g., temperature = 22°C) into a fuzzy set is called:
(A) Defuzzification
(B) Fuzzification
(C) Inference
(D) Aggregation

70. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : A multilayer perceptron (MLP) with a sufficient number of hidden neurons can approximate any continuous function.
Reason (R) : The non-linear activation functions in the hidden layers are crucial for the MLP's ability to model non-linear relationships.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

71. Which of the following is NOT a fundamental component of a communicative act in multi-agent systems?
(A) The sender of the message.

Click Here to Buy book on Amazon

(B) adjust the weights in a neural network to minimize error.
(C) parse natural language sentences.
(D) play games like chess.

127. In a game tree, what is the purpose of the evaluation function used at the cutoff depth?
(A) To determine if the state is a terminal (win/loss/draw) state.
(B) To estimate the expected utility of the game from that state.
(C) To calculate the number of moves made so far.
(D) To check if alpha-beta pruning can be applied.

128. What is the fundamental idea behind Partial Order Planning?
(A) To create a strictly linear sequence of actions.
(B) To delay making ordering decisions until forced to, maintaining flexibility.
(C) To plan in a hierarchical manner from abstract to concrete.
(D) To use a stack to manage goals and subgoals.

129. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : In a multi-agent system, agents need a common language and protocol to communicate effectively.
Reason (R) : Without a shared understanding of message syntax and semantics, coordination and cooperation would be impossible.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

130. What is the key property of a fuzzy set that distinguishes it from a classical (crisp) set?
(A) It can have an infinite number of elements.
(B) Its elements are ordered.
(C) It uses a membership function that assigns a degree of membership between 0 and 1 to elements.
(D) It cannot be empty.

131. In Genetic Algorithms, 'reproduction' involves which of the following main operators?
(A) Selection, Crossover, and Mutation
(B) Encoding and Decoding
(C) Fuzzification and Defuzzification
(D) Initialization and Evaluation

132. The process of breaking down a sentence into its constituent parts (like noun phrase, verb phrase) is called:
(A) Syntactic Analysis (Parsing)
(B) Semantic Analysis
(C) Pragmatic Analysis
(D) Lexical Analysis

133. Which of the following is NOT a component of a classical planning system definition?
(A) A set of states
(B) A set of actions
(C) A goal condition
(D) A fuzzy inference engine

134. Which of the following best describes the 'communicative aspect' of multi-agent systems?
(A) The internal reasoning process of a single agent.
(B) The protocols and languages agents use to interact.
(C) The physical network layer connecting the agents.
(D) The user interface for human operators.

135. What is the output of the 'fuzzification' stage in a fuzzy logic system?
(A) A single crisp numerical value.
(B) A set of fuzzy sets with corresponding degrees of membership.
(C) A boolean true/false value.
(D) The final control action.

136. The 'working principle' of a Genetic Algorithm involves evolving a population of:
(A) candidate solutions to a problem.
(B) logical rules.
(C) neural networks.
(D) state-space representations.

137. Which of the following is a type of unsupervised neural network?
(A) Multilayer Perceptron
(B) Backpropagation Network

Downloaded from www.pdf.tube

193. The property of a fuzzy set is determined by its:
(A) membership function.
(B) number of elements.
(C) name.
(D) logical truth value.

194. In a Genetic Algorithm, 'elitism' is a common strategy where:
(A) the best individual(s) from the current generation are guaranteed to be carried over to the next.
(B) all individuals are mutated in every generation.
(C) only the worst individuals are selected for crossover.
(D) the population size is increased in every generation.

195. Consider the following statements:
Statement I : The Backpropagation algorithm can get stuck in local minima of the error surface.
Statement II : The error surface for a multi-layer neural network is generally non-convex, containing many local minima.
Which of the following is correct?
(A) Both Statement I and Statement II are true, and Statement II is the correct explanation of Statement I.
(B) Both Statement I and Statement II are true, but Statement II is not the correct explanation of Statement I.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

196. The Turing Test falls under which quadrant of the AI approaches matrix?
(A) Thinking Humanly
(B) Acting Humanly
(C) Thinking Rationally
(D) Acting Rationally

197. Which of the following is NOT a part of the STRIPS representation of an action?
(A) Preconditions
(B) Add list
(C) Heuristic value
(D) Delete list

198. In a multi-agent system, the interaction between several independent ride-hailing services competing for customers in a city would be classified as:
(A) Competitive
(B) Cooperative
(C) Hierarchical
(D) Singular

199. Which component of a Genetic Algorithm evaluates how close a given solution is to the optimum?
(A) Fitness function
(B) Encoding
(C) Crossover
(D) Mutation

200. The learning process of a Self-Organizing Map (SOM) is:
(A) Supervised
(B) Unsupervised
(C) Reinforcement-based
(D) Based on back propagation

201. Which knowledge representation method is organized into slots that can be filled with values or pointers to other related structures?
(A) Semantic Networks
(B) Rules
(C) Frames
(D) Scripts

202. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : Alpha-Beta pruning does not affect the final move chosen by a Min-Max search.
Reason (R) : Alpha-Beta pruning is an optimization that only prunes subtrees that it can prove will not influence the final decision at the root node.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

203. The primary purpose of parsing in Natural Language Processing is to:
(A) determine the meaning of a sentence.

NODIA

true, but Statement II is not the correct explanation of Statement I.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

146. Which of the following is NOT a typical component of an agent architecture?
(A) A sensor interface to perceive the environment.
(B) A universal Turing machine for general computation.
(C) An actuator interface to act on the environment.
(D) A program that maps percepts to actions.

147. A cooperative multi-agent system is one in which:
(A) agents work together towards a common objective.
(B) agents have conflicting goals.
(C) agents are unaware of each other's existence.
(D) there is only one agent in the system.

148. What is the role of the 'crossover' operator in Genetic Algorithms?
(A) To introduce small, random changes in a chromosome.
(B) To select the fittest individuals for the next generation.
(C) To combine genetic material from two parent chromosomes to create offspring.
(D) To evaluate the fitness of a chromosome.

149. Arrange the following neural network concepts in the chronological order of their typical development:
1. The Perceptron
2. The concept of back propagation for training multi-layer networks
3. The Hopfield Network for associative memory

Choose the correct sequence:
(A) 1, 3, 2
(B) 1, 2, 3
(C) 2, 1, 3
(D) 3, 1, 2

150. A fuzzy rule-based system for an air conditioner might have a rule like IF temperature is HIGH AND humidity is HIGH, THEN fan_speed is FAST. This is an example of:
(A) a fuzzification process.
(B) a fuzzy inference rule.
(C) a defuzzification method.
(D) a genetic algorithm.

151. Which AI approach is concerned with building systems that think on the basis of a formal, logical framework?
(A) The acting humanly approach (Turing Test).
(B) The thinking humanly approach (Cognitive Modeling).
(C) The acting rationally approach (Rational Agent).
(D) The thinking rationally approach (Laws of Thought).

152. In a state space search, if the heuristic function h(n) is admissible, it means that:
(A) h(n) is always greater than the true cost to the goal.
(B) h(n) never overestimates the true cost to the goal.
(C) h(n) is always equal to the true cost to the goal.
(D) h(n) is computationally inexpensive to calculate.

153. What is a slot in the context of Frame-based knowledge representation?
(A) The entire frame structure.
(B) A procedure that is activated when a value is accessed or changed.
(C) The value that fills an attribute.
(D) A specific attribute or property of the concept being represented.

154. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : Scripts are a useful knowledge representation for understanding stories about stereotyped situations.
Reason (R) : Scripts provide a pre-packaged structure of expected events, roles, and props, which helps in filling in unstated information.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.

Click Here to Buy book on Amazon

(D) A is false but R is true.

155. Which of the following is a key advantage of hierarchical planning over non-hierarchical planning?
(A) It always finds the optimal solution for every problem.
(B) It can solve complex problems with conflicting subgoals perfectly.
(C) It reduces complexity by decomposing the problem into hierarchical levels.
(D) It does not require a formal representation of individual actions.

156. In Natural Language Processing, which phase is responsible for identifying individual words and punctuation, a process also known as tokenization?
(A) Morphological Analysis
(B) Syntactic Analysis
(C) Semantic Analysis
(D) Lexical Analysis

157. Identify the incorrect statement about agent architectures.
(A) A simple reflex agent acts only on the basis of the current percept.
(B) A model-based reflex agent maintains an internal state to track the world.
(C) A goal-based agent has information about a desirable state.
(D) A utility-based agent cannot make decisions when goals are conflicting.

158. Match the items in List I with the items in List II and select the correct option.
List I: (Neural Network Type)
P. Perceptron
Q. Multilayer Perceptron (MLP)
R. Hopfield Network
S. Self-organizing Map (SOM)
List II: (Characteristic)
1. Content-addressable memory
2. Unsupervised clustering and dimensionality reduction
3. Linear classifier
4. Universal function approximator (with hidden layers)
(A) P-3, Q-4, R-1, S-2
(B) P-1, Q-2, R-3, S-4
(C) P-3, Q-1, R-4, S-2
(D) P-4, Q-3, R-2, S-1

159. The 'mutation' operator in a genetic algorithm serves a purpose most similar to:
(A) performing a local search around a good solution.
(B) combining two good solutions.
(C) randomly exploring new parts of the search space.
(D) selecting the best solutions to keep.

160. Which of the following is the most accurate description of the relationship between agents and objects?
(A) Objects are a specific type of agent with added encapsulation.
(B) Agents are a specific type of object with added autonomy.
(C) Agents and objects are completely unrelated concepts.
(D) An agent is a collection of objects, but an object cannot contain an agent.

NODIA

161. Consider the following statements about Alpha-Beta Pruning:
(a) It prunes away branches of the search tree that cannot influence the final decision.
(b) It may change the result of the Min-Max algorithm.
(c) Its efficiency is highly dependent on move ordering.
(d) Beta is the best value for the MAX player found so far on the path.
Which of the above statement(s) is/are correct?
(A) A and B only
(B) A and C only
(C) B and D only
(D) A, C, and D

162. In a fuzzy logic system, what is the role of the inference engine?
(A) To apply the fuzzy rules to the fuzzy inputs to produce a fuzzy output.
(B) To convert crisp inputs to fuzzy values.
(C) To convert fuzzy outputs to crisp values.
(D) To define the membership functions.

163. Which of the following is a primary component of the working principle of a Genetic Algorithm?
(A) The use of gradient information to find a minimum.

(B) The maintenance of a population of candidate solutions.
(C) The formal proof of theorems from axioms.
(D) The construction of a complete state-space graph.

164. What does the back propagation algorithm calculate?
(A) The activation of the output neurons.
(B) The number of layers needed for a network.
(C) The gradient of the error function with respect to the network's weights.
(D) The optimal initial weights for the network.

165. Hopfield networks are a type of recurrent neural network primarily used for:
(A) time-series prediction.
(B) supervised classification of new data.
(C) auto-association and optimization problems.
(D) generating natural language text.

166. A key difference between a model-based reflex agent and a simple reflex agent is the presence of:
(A) sensors.
(B) actuators.
(C) an internal state that represents a model of the world.
(D) a single condition-action rule.

167. Which knowledge representation structure would be most appropriate for capturing the default information that birds can fly, while allowing for exceptions like penguins are birds but cannot fly?
(A) Scripts
(B) Propositional Logic
(C) Frames with inheritance
(D) Conceptual Dependency

168. In Partial Order Planning, a causal link A -> P -> B represents the fact that:
(A) Action A achieves a precondition P required by Action B.
(B) Action A must happen immediately before Action B.
(C) Action A threatens a precondition P required by Action B.
(D) Actions A and B are mutually exclusive.

169. What is the main characteristic of a competitive multi-agent system?
(A) Agents share all information freely.
(B) The goals of the agents are in conflict.
(C) The system is designed to achieve one global, unified goal.
(D) Agents cannot communicate with each other.

170. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : The selection operator in Genetic Algorithms is designed to favor fitter individuals.
Reason (R) : By giving fitter individuals a higher probability of being selected for reproduction, the GA tends to move the population towards better solutions over time.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

171. A Multilayer Perceptron (MLP) is a type of:
(A) feed forward neural network.
(B) recurrent neural network.
(C) unsupervised learning model.
(D) genetic algorithm.

172. The process of identifying word stems (e.g., running -> run) is a part of which Natural Language Processing phase?
(A) Morphological Analysis
(B) Pragmatic Analysis
(C) Semantic Analysis
(D) Syntactic Analysis

173. Identify the incorrect pairing of a concept and its description.
(A) Turing Test - A test for machine intelligence based on conversational abilities.
(B) Rational Agent - An agent that acts to maximize its performance measure.
(C) State Space - The set of all possible states reachable from an initial state.
(D) Heuristic Search - A search that is guaranteed to be optimal but not complete.

Downloaded from www.pdf.tube

(C) Hopfield Network
(D) Radial Basis Function Network (when trained supervised)

138. The primary difference between Supervised and Reinforcement learning is:
(A) Supervised learning uses neural networks, while reinforcement learning does not.
(B) Supervised learning has labeled data with correct outputs, while reinforcement learning uses a reward signal.
(C) Supervised learning is for classification, while reinforcement learning is for regression.
(D) Supervised learning is online, while reinforcement learning is offline.

139. Match the items in List I with the items in List II and select the correct option.
List I: (Genetic Algorithm Term)
P. Chromosome
Q. Gene
R. Population
S. Fitness
List II: (Definition)
1. A measure of how good a solution is
2. A set of individuals (chromosomes) at a given time
3. A part of a chromosome, representing a parameter
4. An encoded representation of a single solution
(A) P-4, Q-3, R-2, S-1
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-2, R-3, S-1
(D) P-2, Q-3, R-4, S-1

140. Which of the following problems is a classic example that cannot be solved by a single perceptron but can be solved by a multi-layer perceptron?
(A) AND gate
(B) OR gate
(C) XOR gate
(D) NOT gate

141. Consider the following statements about planning:
(a) Non-linear planning addresses the problem of interleaved subgoals.
(b) Goal Stack Planning uses a last-in, first-out strategy for handling goals.
(c) Hierarchical planning works by refining high-level goals into low-level actions.
(d) The STRIPS representation uses preconditions, an add list, and a delete list for operators.
(E) All planning algorithms are guaranteed to find the optimal plan.
Which of the above statement(s) is/are correct?
(A) A, B, and E only
(B) C, D, and E only
(C) A, B, C, and D only
(D) B, C, D, and E only

142. The Turing Test is a test of a machine's ability to:
(A) think rationally.
(B) exhibit human-like intelligence in conversation.
(C) act rationally.
(D) perform complex calculations faster than a human.

143. In a game tree, the 'alpha' value associated with a MAX node represents:
(A) the guaranteed minimum score for the MAX player on that path so far.
(B) the guaranteed maximum score for the MIN player on that path so far.
(C) the depth of the node in the tree.
(D) the average score of all its children.

NODIA

144. An ontology is best described as:
(A) a formal specification of the concepts and relationships in a domain.
(B) a collection of rules for an expert system.
(C) a method for handling uncertainty.
(D) a natural language parsing technique.

145. Consider the following statements:
Statement I : Partial Order Planning is often more efficient than Total Order (Linear) Planning for problems with many independent subgoals.
Statement II : This is because Partial Order Planning avoids committing to a specific ordering of actions until a conflict forces it to.
Which of the following is correct?
(A) Both Statement I and Statement II are true, and Statement II is the correct explanation of Statement I.
(B) Both Statement I and Statement II are

P. Sussman Anomaly
Q. Hierarchical Planning
R. Causal Link
S. STRIPS operator

List II: (Description)

1. Decomposing a problem into levels of abstraction
2. Precondition-Add-Delete list formalism
3. A classic problem where a linear planner fails by undoing a needed subgoal
4. A construct in Partial Order Planning where one action enables a precondition for another

(A) P-1, Q-3, R-2, S-4
(B) P-3, Q-1, R-4, S-2
(C) P-3, Q-4, R-1, S-2
(D) P-2, Q-1, R-4, S-3

185. In a multi-agent system, what does autonomy mean?
(A) The agent can operate without direct intervention from humans or other agents.
(B) The agent has complete knowledge of the environment.
(C) The agent is guaranteed to achieve its goals.
(D) The agent is written in a specific programming language.

186. The process of converting a problem's solution representation into a format suitable for a GA, like a bit-string, is called:
(A) Fitness evaluation
(B) Selection
(C) Mutation
(D) Encoding

187. Which of the following is true about a Hop field Network?
(A) It is a feed forward network.
(B) It is a recurrent network that can act as an associative memory.
(C) It is trained using back propagation.
(D) It is an example of unsupervised competitive learning.

188. Learning to play chess by receiving a win or loss signal only at the end of the game is a classic example of:
(A) Reinforcement Learning
(B) Supervised Learning
(C) Unsupervised Learning
(D) Heuristic Learning

189. Which of the following correctly describes the relationship between the Min-Max algorithm and Alpha-Beta pruning?
(A) They are two completely different algorithms for game playing.
(B) Alpha-Beta pruning can sometimes return a different, better move than Min-Max.
(C) Min-Max is an optimization of the Alpha-Beta pruning algorithm.
(D) Alpha-Beta pruning is an optimization that returns the same result as Min-Max but by examining fewer nodes.

190. Consider the following statements about knowledge representation:
(a) Ontologies provide a shared vocabulary and formal specification of a domain.
(b) Semantic networks represent knowledge as a graph of nodes and labeled arcs.
(c) A script for a 'restaurant' scenario would include roles like 'customer' and 'waiter'.
(d) Conceptual Dependency aims to create language-independent representations of meaning.
(E) A frame system can easily represent procedural knowledge.

Which of the above statement(s) is/are correct?
(A) A, B, and E only
(B) C, D, and E only
(C) A, B, C, and D only
(D) A, B, D, and E only

191. What is the primary limitation of Goal Stack Planning that led to the development of non-linear planners?
(A) It cannot handle problems with more than one goal.
(B) It can get stuck in loops or undo previously achieved goals (e.g., Sussman Anomaly).
(C) It is computationally too expensive.
(D) It requires the entire state space to be represented in memory.

192. An agent architecture that maps percepts to actions is the core definition of:
(A) An agent program.
(B) An agent's goal.
(C) An agent's environment.
(D) An agent's utility function.

Click Here to Buy book on Amazon

174. Which of the following is a key feature of a Self-Organizing Map (SOM)?
(A) It learns to map high-dimensional input data onto a lower-dimensional grid.
(B) It requires labeled data for training.
(C) It uses the backpropagation algorithm.
(D) It functions as a content-addressable memory.

175. Consider the following statements regarding AI concepts:
(a) A fuzzy set is defined by a membership function $\mu A(x)$ that can take any real value between 0 and 1.
(b) The crossover operator in a GA combines information from two parent solutions.
(c) A simple reflex agent bases its decisions solely on the current percept, ignoring history.
(d) The Min-Max algorithm is used for single-agent search problems.
Which of the above statement(s) is/are correct?
(A) A and B only
(B) B, C, and D only
(C) A, C, and D only
(D) A, B, and D only

176. Which of the following planning methods explicitly deals with threats as part of its algorithm?
(A) Goal Stack Planning
(B) Hierarchical Planning
(C) Partial Order Planning
(D) STRIPS

177. The communicative aspect of a multi-agent system is primarily concerned with:
(A) how an agent perceives its environment.
(B) the language, protocol, and content of messages exchanged between agents.
(C) the internal data structures of an agent.
(D) the physical hardware that connects the agents.

178. What is the final step in a fuzzy logic controller's process, where a fuzzy set is converted back into a specific, crisp control action?
(A) Defuzzification
(B) Fuzzification
(C) Rule Evaluation
(D) Aggregation

179. In a Genetic Algorithm, if the mutation rate is set too high, what is the likely consequence?
(A) The algorithm's search will become essentially random, and it may not converge.
(B) The algorithm will converge too quickly to a local optimum.
(C) The genetic diversity of the population will decrease rapidly.
(D) The crossover operator will become ineffective.

180. The Back propagation algorithm is a form of:
(A) gradient descent applied to neural networks.
(B) unsupervised clustering.
(C) reinforcement learning.
(D) genetic programming.

181. Which AI approach focuses on creating systems that operate optimally, not necessarily in a human-like way, by selecting actions to maximize a performance measure?
(A) Thinking Humanly
(B) Acting Humanly
(C) Thinking Rationally
(D) Acting Rationally (Rational Agent)

NODIA

182. A key component of an expert system, the 'Inference Engine', does what?
(A) Stores the facts and rules of the domain.
(B) Applies the rules to the known facts to deduce new information.
(C) Provides a way for the user to interact with the system.
(D) Explains the reasoning process to the user.

183. In the context of Natural Language Processing, what is the difference between semantics and pragmatics?
(A) Semantics is about sentence structure, while pragmatics is about meaning.
(B) Semantics is about the literal meaning of words/sentences, while pragmatics is about meaning in context.
(C) Semantics applies to written language, while pragmatics applies to spoken language.
(D) There is no difference; the terms are interchangeable.

184. Match the items in List I with the items in List II and select the correct option.
List I: (Planning Concept)

(B) determine the grammatical structure of a sentence.
(C) translate a sentence to another language.
(D) correct spelling errors in a sentence.

204. Which of the following statements about multi-agent systems is incorrect?
(A) An agent is autonomous.
(B) An agent is situated in an environment.
(C) A multi-agent system consists of a single agent solving a complex task.
(D) Agents in a system can be cooperative or competitive.

205. The 'support' of a fuzzy set is the set of all elements in the universe of discourse that have a membership value:
(A) greater than 0.
(B) equal to 1.
(C) greater than 0.5.
(D) equal to 0.

206. In Genetic Algorithms, the 'crossover' operator is a form of:
(A) exploration.
(B) exploitation.
(C) both exploration and exploitation, by combining existing good solutions.
(D) fitness evaluation.

207. A multilayer perceptron (MLP) consists of an input layer, an output layer, and one or more:
(A) recurrent layers.
(B) hidden layers.
(C) convolutional layers.
(D) memory cells.

208. The key difference between supervised and unsupervised learning lies in the:
(A) availability of labeled training data.
(B) type of neural network used.
(C) complexity of the problem.
(D) use of a reward signal.

209. Arrange the following components of a planning system in a logical hierarchy from most abstract to most specific:
1. Concrete action (e.g., Move(BlockA, Table))
2. Abstract plan (e.g., Clear Block B, then Move Block B)
3. High-level goal (e.g., On(B, C))

Choose the correct sequence:
(A) 3, 2, 1
(B) 1, 2, 3
(C) 2, 3, 1
(D) 3, 1, 2

210. The concept of an 'agent' being autonomous, proactive, reactive, and social is central to the field of:
(A) Knowledge Representation.
(B) Multi-Agent Systems.
(C) Natural Language Processing.
(D) Fuzzy Logic.

211. Which of the following is NOT a valid approach to AI?
(A) Turing Test approach
(B) Emotional Intelligence approach
(C) Rational Agent approach
(D) State Space Representation approach

212. Match the items in List I with the items in List II and select the correct option.
List I: (Term)
P. Ontology
Q. STRIPS
R. Back propagation
S. Alpha-Beta
List II: (Concept)
1. Game Tree Pruning
2. Neural Network Training
3. Planning Operator Representation
4. Formal Knowledge Specification
(A) P-4, Q-3, R-2, S-1
(B) P-1, Q-2, R-3, S-4
(C) P-4, Q-2, R-3, S-1
(D) P-2, Q-3, R-4, S-1

NODIA

213. Handling uncertainty in knowledge is a key problem addressed by all of the following EXCEPT:
(A) Fuzzy Logic
(B) Bayesian Networks
(C) Certainty Factors
(D) First-Order Predicate Logic

214. Which of the following statements about planning is true?
(A) Non-linear planning addresses subgoal interactions by interleaving plans.
(B) Linear planning creates a partially ordered

plan.
(C) Goal Stack Planning is an example of a non-linear planner.
(D) Hierarchical planning only works on single-level problems.

215. Consider the following statements about Genetic Algorithms:
(a) The 'encoding' step defines how a solution is represented.
(b) The 'fitness function' is problem-specific.
(c) 'Crossover' exchanges information between two solutions.
(d) 'Mutation' helps prevent premature convergence to local optima.
(E) The algorithm operates on a single candidate solution at a time.
Which of the above statement(s) is/are correct?
(A) A, B, C, and D only
(B) A, B, C, and E only
(C) B, C, D, and E only
(D) A, C, D, and E only

216. A Hopfield network's state converges to a local minimum of an energy function. This makes it useful as a:
(A) classifier for unseen data.
(B) generator of novel sequences.
(C) content-addressable memory.
(D) tool for feature extraction.

217. In a rational agent, the choice of action depends on the:
(A) current percept only.
(B) percept sequence, knowledge, and goals.
(C) physical architecture of the agent only.
(D) agent's programming language.

218. The problem of determining the structure of a sentence (e.g., identifying noun phrases, verb phrases) is called:
(A) Pragmatics
(B) Semantics
(C) Parsing
(D) Tokenization

219. The Sussman Anomaly highlights a fundamental issue with which type of planning?
(A) Hierarchical Planning
(B) Partial-Order Planning
(C) Simple Linear Planning (like Goal Stack Planning)
(D) Case-Based Planning

220. The following question consists of two statements, an Assertion (A) and a Reason (R). Select the correct answer using the code given below.
Assertion (A) : Self-Organizing Maps (SOMs) are a form of unsupervised learning.
Reason (R) : SOMs learn to create a discretized representation of the input space (a map) without needing labeled data, identifying patterns and similarities on their own.
(A) Both A and R are true and R is the correct explanation of A.
(B) Both A and R are true but R is not the correct explanation of A.
(C) A is true but R is false.
(D) A is false but R is true.

221. Which knowledge representation scheme uses nodes to represent objects/concepts and arcs to represent relations between them in a graph-like structure?
(A) Frames
(B) Logic
(C) Scripts
(D) Semantic Networks

222. What is the role of an 'ontology' in knowledge representation?
(A) To provide a set of inference rules.
(B) To handle uncertainty and probability.
(C) To create a formal, explicit specification of the concepts and relationships in a domain.
(D) To represent a sequence of events.

223. In a multi-agent system where agents from different organizations must bid for resources, their interaction is best described as:
(A) Competitive
(B) Cooperative
(C) Reflexive
(D) Passive

224. Which of the following is a core component of a fuzzy rule-based system?
(A) A set of crisp if-then rules
(B) A set of fuzzy if-then rules and membership functions
(C) A back propagation algorithm

Downloaded from www.pdf.tube

(D) A state-space search graph

225. Consider the following statements about Artificial Neural Networks:
(a) A single perceptron can learn the XOR function.
(b) Back propagation is an algorithm for supervised learning.
(c) A Hopfield network is a form of recurrent network.
(d) A Self-Organizing Map produces a low-dimensional representation of input data.
Which of the above statement(s) is/are correct?
(A) A and B only
(B) B, C, and D only
(C) A, C, and D only
(D) A, B, and D only

226. The `fitness function` in a Genetic Algorithm is crucial because it:
(A) generates the initial population.
(B) performs the crossover operation.
(C) guides the search by evaluating the quality of solutions.
(D) introduces random variation into the population.

227. In the context of the back propagation algorithm, the error is propagated from the:
(A) output layer back to the input layer.
(B) input layer to the output layer.
(C) hidden layers to the input layer.
(D) first hidden layer to the last hidden layer.

228. Which of these tasks is an example of Reinforcement Learning?
(A) Classifying images of cats and dogs from a labeled dataset.
(B) Grouping customers into different market segments based on their purchasing history.
(C) Training a robot to navigate a maze by rewarding it for moving closer to the exit.
(D) Predicting house prices based on historical sales data.

229. A key assumption of the Min-Max algorithm is that the opponent:
(A) plays randomly.
(B) plays sub-optimally.
(C) plays optimally to maximize their own score.
(D) will always choose the move that leads to the quickest game end.

230. The Conceptual Dependency framework represents knowledge in terms of:
(A) a graph of interconnected concepts.
(B) a small set of conceptual primitives representing actions, states, and objects.
(C) a hierarchy of frames and slots.
(D) a set of if-then rules.

NODIA

231. The primary purpose of Hierarchical Planning is to:
(A) ensure the plan is always optimal.
(B) manage complexity by planning at different levels of abstraction.
(C) handle interacting subgoals perfectly.
(D) produce a strictly linear plan.

232. In Natural Language Processing, the study of how language is used in context to achieve goals (e.g., making requests, giving commands) is known as:
(A) Syntax
(B) Semantics
(C) Pragmatics
(D) Morphology

233. An 'agent architecture' can be described as a blueprint for an agent that maps:
(A) percepts to actions.
(B) actions to goals.
(C) goals to utilities.
(D) percepts to states.

234. In fuzzy logic, the operation for the union of two fuzzy sets A and B ($A \cup B$) is usually defined by the:
(A) minimum of their membership functions.
(B) maximum of their membership functions.
(C) product of their membership functions.
(D) average of their membership functions.

235. The three fundamental genetic operators in a simple Genetic Algorithm are:
(A) Selection, Crossover, Mutation
(B) Encoding, Fitness, Decoding
(C) Initialization, Evaluation, Termination
(D) Perceptron, Backpropagation, Hopfield

(B) It is an explicit and formal specification of a conceptualization.
(C) It is primarily used for representing sequences of events.
(D) It is dependent on the natural language used to create it.

279. What key function of ontologies is highlighted in the text?
(A) To ensure that AI systems pass the Turing Test.
(B) To break down sentences into primitive actions and conceptual roles.
(C) To enable knowledge sharing and interoperability between different systems.
(D) To find the optimal move in a two-player adversarial game.

280. The passage implies that a significant difference between an ontology and a simpler semantic network is the ontology's:
(A) use of nodes and arcs.
(B) ability to represent 'is-a' relationships.
(C) focus on natural language understanding.
(D) higher degree of formality and use of standard languages.

Direction : Please read the following source carefully.

An expert system is a sophisticated AI application engineered to emulate the decision-making ability of a human expert within a specific, narrow domain. The fundamental architecture of such a system is bifurcated into two core components: the knowledge base and the inference engine. The knowledge base serves as the repository for domain-specific knowledge, which is meticulously extracted from human experts and encoded in a machine-readable format, often as a collection of facts and 'IF-THEN' rules. The inference engine, in contrast, is the reasoning component of the system. It is a general-purpose mechanism that systematically applies the rules from the knowledge base to the facts of a given problem to deduce new conclusions or recommend actions. A critical challenge for real-world expert systems is that the information they work with is often incomplete, ambiguous, or inherently uncertain. Therefore, a robust expert system must have a mechanism for handling uncertainty. Early systems used methods like certainty factors, which associate a numerical degree of belief with facts and rules. More modern approaches often employ probabilistic reasoning models, such as Bayesian networks, which allow the system to manage and propagate uncertainty according to the principles of probability theory, resulting in conclusions that are qualified with a degree of confidence.

Based on the above source, attempt the following five questions.

281. What are the two primary components of an expert system's architecture as described in the text?
(A) A knowledge base and an inference engine.
(B) A user interface and a database.
(C) A set of sensors and a set of actuators.
(D) A parser and a semantic analyzer.

282. What is the specific function of the inference engine in an expert system?
(A) To store the domain-specific facts and rules.
(B) To apply knowledge to a problem to derive conclusions.
(C) To acquire and encode knowledge from human experts.
(D) To provide a user-friendly interface for interaction.

283. According to the passage, why is handling uncertainty a crucial capability for expert systems?
(A) Because the inference engine itself is an uncertain process.
(B) Because real-world problem domains often involve incomplete or uncertain information.
(C) Because knowledge from human experts is always contradictory.
(D) Because the knowledge base has a limited storage capacity.

284. The passage mentions certainty factors and Bayesian networks as examples of what?
(A) Techniques for managing and reasoning with uncertain information.
(B) Methods for representing knowledge in the knowledge base.
(C) Algorithms used by the inference engine for logical deduction.
(D) Architectures for structuring the entire expert system.

Click Here to Buy book on Amazon

(B) Name, Parameters, and Body.
(C) Cost, Heuristic, and Goal.
(D) Input, Output, and State.

249. Which of the following is a classic example of an unsupervised learning neural network?
(A) Multilayer Perceptron
(B) Backpropagation Network
(C) Self-Organizing Map (SOM)
(D) Perceptron

250. The core idea behind the Turing Test is to define intelligence in terms of:
(A) computational power.
(B) logical perfection.
(C) rational decision making.
(D) conversational indistinguishability from a human.

Direction : Please read the following source carefully.

The field of Artificial Intelligence (AI) encompasses diverse philosophical and methodological approaches to creating intelligent systems. A seminal concept is the Turing Test, proposed by Alan Turing, which defines intelligence in terms of behavioral equivalence. An AI is said to pass the test if a human interrogator, after posing written questions, cannot reliably distinguish the machine's responses from those of a human. This approach focuses on the ability to exhibit human-like intelligent behavior, or thinking humanly. In contrast, the rational agent approach provides a more formal and scientifically measurable framework. An agent is conceptualized as any entity that perceives its environment through sensors and acts upon that environment through actuators. A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome, based on its percept sequence and built-in knowledge. This paradigm shifts the focus from mimicking human thought to achieving optimal performance against a defined measure. The emphasis on rationality and performance makes this approach more tractable for engineering and scientific analysis, as rationality can be mathematically defined and optimized, whereas the criteria for thinking humanly remain more ambiguous and subjective.

Based on the above source, attempt the following five questions.

NODIA

251. What is the primary criterion for success in the Turing Test?
(A) The machine's ability to solve complex mathematical problems.
(B) The machine's indistinguishability from a human in a textual conversation.
(C) The machine's capacity to perceive its environment through sensors.
(D) The machine's efficiency in achieving a pre-defined performance measure.

252. According to the passage, what defines a rational agent?
(A) An entity that acts to maximize its performance measure based on available information.
(B) An entity that perfectly mimics human cognitive processes.
(C) An entity that passes the Turing Test under all conditions.
(D) An entity that possesses actuators but not necessarily sensors.

253. The passage contrasts the Turing Test with the rational agent approach. What is a key reason cited for the rational agent approach being more scientifically tractable?
(A) It requires less computational power to implement.
(B) It focuses exclusively on natural language conversation.
(C) It is based on a well-defined and measurable concept of rationality.
(D) It was developed more recently than the Turing Test.

254. Which of the following best describes the components of an agent as defined in the source?
(A) A knowledge base and an inference engine.
(B) A set of rules and logical propositions.
(C) A performance measure and a utility function.
(D) A mechanism for perception (sensors) and a mechanism for action (actuators).

255. Based on the text, an AI system designed to play chess by making the moves most likely to lead to a win, rather than moves a human grandmaster would make, would be best described as adhering to which approach?
(A) The Turing Test approach

or objects and directed arcs represent the relationships between them. A key feature of semantic networks is the use of specific link types, such as the 'is-a' link, which denotes a subclass relationship. This structure facilitates a type of reasoning known as inheritance, where properties of a general class node are automatically inherited by more specific subclass nodes.

Based on the above source, attempt the following five questions.

266. What is the key advantage of First-Order Logic (FOL) over Propositional Logic as described in the passage?
(A) FOL is computationally less expensive than Propositional Logic.
(B) FOL can represent relationships between objects and make generalized statements.
(C) FOL uses a simpler syntax with fewer logical operators.
(D) FOL is a graphical representation that is easier to visualize.

267. In a semantic network, what is the typical function of the nodes and arcs?
(A) Nodes represent logical operators, and arcs represent propositions.
(B) Nodes represent concepts, and arcs represent relationships between them.
(C) Nodes represent actions, and arcs represent preconditions.
(D) Nodes represent quantifiers, and arcs represent variables.

268. The ability of a 'dog' node to inherit the property 'has a tail' from an 'animal' node in a semantic network is an example of what process?
(A) Inheritance
(B) Quantification
(C) Logical deduction
(D) Adversarial search

269. According to the text, what construct in First-Order Logic allows for making statements about all or some objects in a domain?
(A) Predicates
(B) Variables
(C) Quantifiers
(D) Logical operators

270. Based on the passage, which representation scheme is inherently graphical in nature?
(A) Propositional Logic
(B) First-Order Logic
(C) Rational Agent Model
(D) Semantic Network

Direction : Please read the following source carefully.

To organize complex knowledge beyond simple logical statements, AI utilizes structured representation schemes like frames, rules, and scripts. A frame is a data structure designed to represent a stereotyped situation, such as the features of a room or the details of an event. It contains 'slots' that hold specific information pertinent to the situation. For instance, a 'room' frame might have slots for 'color', 'flooring', and 'contents'. Slots can be filled with values, and importantly, they can hold default values that apply in the absence of more specific information. This allows for assumption-based reasoning. Another powerful scheme is the rule-based system, which captures knowledge in the form of 'IF-THEN' conditional statements. The 'IF' part, or antecedent, specifies a condition, and the 'THEN' part, or consequent, specifies a conclusion or action. When the facts in a knowledge base satisfy the condition of a rule, the rule is said to 'fire', and its conclusion is added to the knowledge base. For representing procedural knowledge, scripts are employed. A script is a specialized structure that outlines a stereotyped sequence of events for a particular context, like visiting a restaurant. It contains elements such as entry conditions, results, props, roles, and a sequence of scenes, enabling an AI to make inferences about the typical flow of events.

Based on the above source, attempt the following five questions.

271. In a frame-based representation, what is the purpose of a 'slot'?
(A) To define the sequence of events in a situation.
(B) To represent an 'IF-THEN' conditional statement.
(C) To hold a specific piece of information or attribute related to the concept.
(D) To connect one frame to another in a hierarchy.

Downloaded from www.pdf.tube

performs a depth-first exploration of the game tree, assigning a utility value to each terminal state (e.g., win, lose, draw). These values are then propagated up the tree; MAX nodes select the child with the maximum value, and MIN nodes select the child with the minimum value. A major drawback of the standard Min-Max algorithm is its computational complexity, as it must explore the entire game tree up to a specified depth. To address this inefficiency, the Alpha-Beta cutoff procedure, or pruning, was developed. This is an optimization that can be applied to the Min-Max tree search. It works by pruning large parts of the search tree that cannot possibly influence the final decision. It maintains two values during the search: alpha, representing the best score found so far for MAX, and beta, representing the best score found so far for MIN. If a path is found to be worse than a previously examined path for either player, that branch of the tree is ignored.

Based on the above source, attempt the following five questions.

261. What is the fundamental assumption underlying the Min-Max algorithm?
(A) The opponent will make random moves.
(B) The game will always end in a win or a loss, but not a draw.
(C) The game tree is always shallow enough to be fully explored.
(D) The opponent will also play optimally to achieve their best outcome.

262. In the context of the Min-Max algorithm, what is the role of the MAX player?
(A) To minimize the final utility value of the game.
(B) To select moves that lead to states with the maximum possible score.
(C) To prune branches of the game tree using the alpha value.
(D) To assign utility values to the terminal states of the game.

263. What is the primary purpose of the Alpha-Beta cutoff procedure?
(A) To increase the depth of the game tree search.
(B) To replace the utility function with a more accurate one.
(C) To reduce the number of nodes evaluated by the Min-Max algorithm.
(D) To ensure that the opponent is playing optimally.

NODIA

264. How does Alpha-Beta pruning achieve its optimization?
(A) By estimating the utility of non-terminal nodes using a heuristic.
(B) By converting the game tree into a simpler graph structure.
(C) By exploring the tree in a breadth-first manner instead of depth-first.
(D) By ignoring branches that are guaranteed not to affect the final optimal decision.

265. The passage states that a major drawback of the standard Min-Max algorithm is its:
(A) High computational cost due to exhaustive exploration.
(B) Inability to handle games with more than two players.
(C) Requirement for the game to have perfect information.
(D) Difficulty in assigning utility values to terminal states.

Direction : Please read the following source carefully.

Knowledge representation and reasoning are cornerstones of artificial intelligence, providing the formalisms to encode and manipulate information. Logic serves as a powerful and foundational tool for this purpose. Propositional logic, the simpler form, deals with propositions—declarative sentences that can be either true or false—and connects them using logical operators like AND, OR, and NOT. While useful, its expressive power is limited as it cannot represent relationships between objects or make generalized statements. First-Order Logic (FOL) overcomes this limitation by introducing predicates, variables, and quantifiers. Predicates describe properties of objects or relations among them, while variables allow for abstraction. Quantifiers, namely the universal quantifier (∀, for all) and the existential quantifier (∃, there exists), enable the expression of general rules and facts about entire collections of objects. An alternative, more graphical approach to knowledge representation is the semantic network. This formalism uses a graph structure where nodes represent concepts

carefully.

Various techniques have been developed to implement AI planning systems. One of the earliest approaches is Goal Stack Planning. This method utilizes a last-in, first-out stack to manage the planning process. The main goal is initially pushed onto the stack. The planner then repeatedly pops the top item. If the item is a compound goal, it is replaced on the stack by its constituent subgoals. If the item is a single, unsatisfied goal, an appropriate action is chosen and pushed onto the stack, followed by its preconditions. If the item is an action, its preconditions are checked; if satisfied, the action is simulated by updating the world state. A foundational representation used by many planners, including early ones, is the STRIPS (Stanford Research Institute Problem Solver) representation for actions. In STRIPS, an operator is defined by three lists: a set of preconditions, an add list (literals that the action makes true), and a delete list (literals that the action makes false). For more complex problems with many interacting subgoals, Hierarchical Planning is often more effective. This strategy involves planning at a high level of abstraction first, creating a simplified plan with abstract operators. This abstract plan is then progressively refined by decomposing the abstract steps into more concrete actions, which tackles the most significant parts of the problem first and reduces the overall search complexity.

Based on the above source, attempt the following five questions.

291. In Goal Stack Planning, what happens when an action is popped from the stack and its preconditions are satisfied?
(A) The action's effects are pushed onto the stack as new goals.
(B) The planner backtracks to a previous state.
(C) The action is simulated, and the world state is updated.
(D) The main goal is pushed back onto the stack for re-evaluation.

292. What are the three components that define an action in the STRIPS representation?
(A) Preconditions, Postconditions, and Effects.
(B) State, Goal, and Operator.
(C) Actor, Object, and Instrument.
(D) Preconditions, Add list, and Delete list.

293. In the STRIPS formalism, what is the function of the 'add list'?
(A) It lists the new facts that become true after the action is executed.
(B) It specifies the conditions that must be true before the action can be run.
(C) It enumerates the facts that are no longer true after the action is executed.
(D) It contains a list of alternative actions to be considered.

294. What is the core principle behind Hierarchical Planning?
(A) To explore all possible action sequences in a breadth-first manner.
(B) To create a plan at an abstract level first and then refine it with details.
(C) To use a stack to manage goals and subgoals in a last-in, first-out order.
(D) To represent all actions using only an add list and a delete list.

295. The passage suggests that Hierarchical Planning is particularly advantageous for what kind of problems?
(A) Problems with simple, non-interacting subgoals.
(B) Problems that can be solved with a short, linear sequence of actions.
(C) Complex problems where abstracting away details can reduce search.
(D) Problems with a very small state space.

Direction : Please read the following source carefully.

Natural Language Processing (NLP) is a subfield of AI dedicated to enabling computers to process and understand human language. A critical component of NLP is grammar, which provides a formal set of rules defining how words can be combined to form syntactically valid sentences in a language. The process of applying these rules to a sentence to determine its structure is known as parsing. A parser is an algorithm that takes an input sentence and attempts to produce a parse tree—a hierarchical structure that graphically represents how the grammar rules can generate the sentence. This process verifies the sentence's syntactic correctness. Parsing techniques are generally categorized into two types: top-down parsing, which begins with the highest-level

Click Here to Buy book on Amazon

272. How is knowledge primarily represented in a rule-based system?
(A) As a network of nodes and arcs.
(B) As a collection of 'IF-THEN' conditional statements.
(C) As a sequence of stereotypical scenes.
(D) As a data structure with slots and default values.

273. Which knowledge representation scheme is specifically designed to model a stereotyped sequence of events?
(A) Frame
(B) Rule
(C) Semantic Network
(D) Script

274. What happens when a rule 'fires' in a rule-based system?
(A) The action or conclusion specified in the rule's 'THEN' part is executed.
(B) The system requests more information from the user.
(C) The condition of the rule is proven to be false.
(D) The default values in a corresponding frame are updated.

275. The concept of having a 'default value' for a 'wall color' slot in a 'bedroom' frame is a feature of which representation?
(A) Scripts, as they define default scenes.
(B) Rules, as they have default conclusions.
(C) Frames, as they allow for assumption-based reasoning.
(D) First-Order Logic, as it uses universal quantifiers.

Direction : Please read the following source carefully.

For deeper semantic understanding, particularly in natural language processing, AI employs advanced knowledge representation schemes like Conceptual Dependency and ontologies. The Conceptual Dependency (CD) theory aims to represent the meaning of sentences in a canonical, language-independent form. This is achieved by breaking down sentences into a structure composed of a small set of primitive actions (e.g., PTRANS for physical location transfer, ATRANS for abstract possession transfer) and their conceptual roles (actor, object, instrument). The primary goal is to ensure that sentences with identical meanings have identical representations, regardless of their surface linguistic structure, thereby simplifying semantic inference. On a broader and more formal scale, an ontology is defined as an explicit, formal specification of a shared conceptualization. It serves as a structured knowledge model that defines a vocabulary of concepts, properties, and relationships within a specific domain. Unlike simpler schemes like semantic networks, ontologies are often constructed using standardized languages like the Web Ontology Language (OWL) and are designed for rigorous, machine-interpretable semantics. Their principal function is to enable knowledge sharing, data integration, and semantic interoperability across disparate AI systems and agents, playing a critical role in the development of the Semantic Web.

Based on the above source, attempt the following five questions.

276. What is the central goal of Conceptual Dependency (CD) theory?
(A) To represent the meaning of sentences in a standardized, language-independent format.
(B) To create a formal specification of a shared conceptualization for a domain.
(C) To parse the grammatical structure of sentences using a set of formal rules.
(D) To model stereotyped sequences of events in specific contexts.

277. In Conceptual Dependency, what is the role of primitive actions like 'PTRANS' and 'ATRANS'?
(A) They represent the grammatical parts of speech in a sentence.
(B) They define the properties and relationships between concepts in an ontology.
(C) They act as variables that can be quantified over a domain of objects.
(D) They serve as the fundamental building blocks for representing the meaning of actions.

278. According to the passage, what is a defining characteristic of an ontology?
(A) It is an informal, intuitive representation of knowledge.

Downloaded from www.pdf.tube

236. The Backpropagation algorithm is fundamentally an application of which mathematical concept to train neural networks?
(A) The chain rule for derivatives
(B) Bayesian probability
(C) Linear algebra matrix operations
(D) Fuzzy set theory

237. A Self-Organizing Map (SOM) is particularly useful for:
(A) supervised classification.
(B) reinforcement learning.
(C) visualization and exploratory data analysis.
(D) symbolic reasoning.

238. Consider the following statements:
Statement I : A heuristic function h(n) is called admissible if it never overestimates the cost to reach the goal.
Statement II : A* search using an admissible heuristic is guaranteed to find an optimal solution.
Which of the following is correct?
(A) Both Statement I and Statement II are true, and Statement II is a direct consequence of Statement I.
(B) Both Statement I and Statement II are true, but are unrelated.
(C) Statement I is true, but Statement II is false.
(D) Statement I is false, but Statement II is true.

239. The knowledge representation that uses scenes and a temporal sequence of events to represent stereotyped situations is a:
(A) Frame
(B) Semantic Network
(C) Script
(D) Rule-based system

240. Partial Order Planning (POP) differs from Total Order Planning (TOP) because:
(A) POP is less flexible.
(B) POP commits to ordering decisions only when necessary to resolve conflicts.
(C) TOP is more efficient for problems with independent subgoals.
(D) POP cannot handle complex goals.

241. Which of the following is NOT a type of learning in Artificial Neural Networks?
(A) Supervised Learning
(B) Unsupervised Learning
(C) Deductive Learning
(D) Reinforcement Learning

242. In a game tree search, the alpha value represents the best score for the _____ player, and the beta value represents the best score for the _____ player.
(A) MAX, MIN
(B) MIN, MAX
(C) current, next
(D) human, computer

NODIA

243. The process of analyzing a sentence to determine its meaning is known as:
(A) Semantic analysis
(B) Syntactic analysis
(C) Pragmatic analysis
(D) Lexical analysis

244. What is the fundamental building block of a simple Genetic Algorithm's representation?
(A) A neuron
(B) A rule
(C) A chromosome
(D) A state

245. A single-layer Perceptron can only solve problems that are:
(A) non-linearly separable.
(B) linearly separable.
(C) related to time-series data.
(D) solvable using unsupervised learning.

246. In a fuzzy rule-based system, the purpose of 'defuzzification' is to:
(A) convert a crisp input to a fuzzy set.
(B) apply fuzzy rules to generate a fuzzy output.
(C) convert the final fuzzy output set into a single crisp value.
(D) define the linguistic variables.

247. The difference between an agent and a standard software object is primarily the agent's:
(A) autonomy and proactive behavior.
(B) encapsulation of data.
(C) use of methods and properties.
(D) ability to be instantiated.

248. The STRIPS planner represents an action by its:
(A) Preconditions, Add list, and Delete list.

285. The knowledge base of an expert system is typically populated with information sourced from:
(A) General-purpose encyclopedias and textbooks.
(B) The system's own trial-and-error learning.
(C) Human experts in a specific, narrow domain.
(D) Large, unstructured datasets from the internet.

Direction : Please read the following source carefully.

In artificial intelligence, planning is the deliberative process of deciding on a sequence of actions to achieve a specific goal. A classical planning system operates on a formalized representation of the world. The essential components of this representation include a description of states, goals, and actions. States are typically represented as a conjunction of positive literals, describing what is true in the world at a particular moment. A goal is a partial description of a desired state, specified as a conjunction of literals to be achieved. Actions, or operators, are the means of changing the world state. Each action is defined by its preconditions, which are the literals that must be true for the action to be executable, and its effects, which describe the changes to the state (literals that become true or false) after the action is performed. Planning algorithms can be classified based on how they order actions. Linear planning, also known as totally ordered planning, constructs a plan as a single, strict sequence of steps. This approach can be inefficient as it commits to an ordering early, which may lead to dead ends and require significant backtracking. In contrast, non-linear planning, also referred to as partial-order planning, is more flexible. It maintains a partial ordering of actions and only adds ordering constraints when necessary to resolve conflicts, allowing subgoals to be worked on independently.

Based on the above source, attempt the following five questions.

286. According to the passage, what is the fundamental purpose of AI planning?
(A) To represent the world as a conjunction of literals.
(B) To classify actions based on their preconditions and effects.
(C) To perceive the environment and act upon it rationally.
(D) To find a sequence of actions that leads from an initial state to a goal state.

287. In a classical planning system, what are the 'preconditions' of an action?
(A) The changes that occur in the world after the action is executed.
(B) The set of goals the action is designed to achieve.
(C) The conditions that must be true for the action to be executable.
(D) The sequence of actions that must precede the current action.

288. What is the key characteristic of a linear planner?
(A) It allows actions in the plan to be unordered relative to each other.
(B) It generates a plan as a strict, sequential ordering of actions.
(C) It represents states using a non-linear data structure.
(D) It avoids backtracking by resolving all conflicts in advance.

289. The passage highlights which of the following as an advantage of non-linear planning over linear planning?
(A) It guarantees finding a solution more quickly.
(B) It requires a simpler representation for actions.
(C) It eliminates the need for defining action preconditions.
(D) It is more flexible as it does not commit to a specific action order unless necessary.

290. How is a 'goal' typically represented in a classical planning system?
(A) As a partial description of a desired state to be achieved.
(B) As a complete description of the final world state.
(C) As a sequence of actions to be performed.
(D) As a heuristic function estimating cost to completion.

Direction : Please read the following source

Click Here to Buy book on Amazon

(B) The thinking humanly approach
(C) The rational agent approach
(D) The behavioral equivalence approach

Direction : Please read the following source carefully.

In artificial intelligence, many problems can be formulated as a search for a solution within a state space. The state space representation formalizes a problem by defining a set of possible states the problem can be in, an initial state where the search begins, and one or more goal states that represent the solution. The transitions between these states are governed by a set of well-defined operators or actions. This framework effectively transforms problem-solving into a graph traversal problem: finding a path from the initial node to a goal node. For problems with a small number of states, uninformed search algorithms like Breadth-First Search (BFS) might suffice. However, for complex domains with vast, combinatorially explosive state spaces, such exhaustive methods become computationally infeasible. This is where heuristic search techniques become essential. A heuristic search is an informed search strategy that uses a heuristic function, denoted as h(n), to guide the search process. This function provides an estimate of the cost to reach a goal from a given state 'n'. By prioritizing states that appear to be closer to the goal, heuristic algorithms can drastically prune the search space and find solutions much more efficiently than their uninformed counterparts. The quality of the heuristic function is paramount; a well-designed heuristic can lead to significant performance gains.

Based on the above source, attempt the following five questions.

256. Which of the following are the essential components of a state space representation as described in the passage?
(A) A heuristic function, an initial state, and a set of operators.
(B) An initial state, a set of operators, and one or more goal states.
(C) A set of nodes, a set of edges, and a heuristic function.
(D) A performance measure, a set of sensors, and a set of actuators.

257. Why are heuristic search techniques often preferred over uninformed search methods for complex problems?
(A) They guarantee finding the shortest possible path to the goal.
(B) They are simpler to implement than algorithms like BFS.
(C) They systematically explore every possible state to ensure completeness.
(D) They use problem-specific knowledge to reduce the search space and improve efficiency.

258. What is the primary role of the heuristic function, h(n), in a heuristic search?
(A) To estimate the cost of reaching a goal from the current state 'n'.
(B) To calculate the exact cost of the path from the initial state to state 'n'.
(C) To determine if a state 'n' is a goal state.
(D) To define the set of legal operators applicable at state 'n'.

259. The passage suggests that the effectiveness of a heuristic search is most dependent on what?
(A) The size of the computer's memory.
(B) The processing speed of the CPU.
(C) The quality of the guiding heuristic function.
(D) The number of goal states in the problem.

260. According to the text, framing a problem using a state space representation allows problem-solving to be viewed as:
(A) A task of finding a path in a graph.
(B) A process of logical deduction.
(C) A challenge of maximizing a performance measure.
(D) A simulation of human cognitive processes.

NODIA

Direction : Please read the following source carefully.

Adversarial search algorithms are fundamental to AI applications involving competitive environments, such as two-player games with perfect information. The cornerstone of these algorithms is the Min-Max search. This technique explores a game tree to determine the optimal move for the current player, typically designated as MAX. The algorithm operates under the assumption that the opposing player, MIN, will also play optimally to counter MAX's moves. It

grammar rule (e.g., S for sentence) and attempts to derive the sentence, and bottom-up parsing, which starts with the individual words of the sentence and tries to build up to the top-level rule. While parsing establishes syntactic validity, it does not determine meaning. This is the objective of semantic analysis. Semantic analysis is the process that follows parsing and focuses on interpreting the meaning of the syntactic structure. It involves mapping words and grammatical structures to a formal meaning representation, resolving ambiguities (e.g., word sense, pronoun reference), and ultimately deriving the intended message conveyed by the text.

Based on the above source, attempt the following five questions.

296. What is the primary function of a grammar in the context of NLP?
(A) To define the meaning of individual words.
(B) To resolve ambiguity in pronoun references.
(C) To map sentences to a language-independent representation.
(D) To provide rules for forming syntactically correct sentences.

297. The process of analyzing a sentence to determine its grammatical structure is called:
(A) Semantic analysis
(B) Tokenization
(C) Parsing
(D) Heuristic search

298. What is the fundamental difference between top-down and bottom-up parsing?
(A) Top-down starts from the sentence symbol, while bottom-up starts from the words.
(B) Top-down is faster, while bottom-up parsing is more accurate.
(C) Top-down produces a parse tree, while bottom-up produces a logical form.
(D) Top-down is used for syntax, while bottom-up is used for semantics.

299. According to the passage, why is semantic analysis a necessary step after parsing?
(A) Because parsing cannot handle grammatically incorrect sentences.
(B) Because syntactic correctness does not guarantee that a sentence is meaningful.
(C) Because parsing does not produce a parse tree.
(D) Because semantic analysis is required to identify the parts of speech.

300. Deriving the intended message of a sentence by resolving ambiguity and interpreting its structure is the goal of which process?
(A) Semantic analysis
(B) State space representation
(C) Syntactic parsing
(D) Hierarchical planning

www.nodia.in

Downloaded from www.pdf.tube

ANSWERS

1. (B) To assess if a machine's behavior is indistinguishable from that of a human.
The Turing Test, proposed by Alan Turing, evaluates a machine's ability to exhibit intelligent behavior equivalent to or indistinguishable from a human. In this imitation game, a human judge engages in natural language conversations with a machine and another human. If the judge cannot reliably tell the machine from the human, the machine passes the intelligence test successfully.

2. (A) An agent that acts to achieve the best possible or best expected outcome.
A rational agent is a conceptual framework in artificial intelligence describing an entity capable of perceiving its environment and taking actions. These agents operate by selecting actions that maximize their performance measure. This approach focuses on achieving optimal results based on available information, rather than strictly mimicking human cognitive processes or emotional responses during the decision-making process.

3. (C) The exact sequence of actions to reach the solution.
A state space representation defines a problem in terms of its initial state, a set of operators for transitioning between states, and a goal condition. While it provides the formal framework for finding a solution, the specific sequence of actions is the result of a search process, rather than a predefined component of the state space structure itself.

4. (C) Statement I is true, but Statement II is false.
Heuristic search techniques utilize problem-specific knowledge, often represented as an evaluation function, to navigate the search space more efficiently than blind search methods. While these techniques significantly improve performance, they do not guarantee finding an optimal solution in every case unless specific conditions, such as admissibility in the A* algorithm, are strictly met by the chosen heuristic function.

5. (C) Find the optimal move assuming the opponent also plays optimally.
The Min-Max algorithm is a decision-making rule used in zero-sum games to determine the best move for a player. It operates on the assumption that the opponent will also play perfectly to minimize the first player's gain. By exploring the game tree, it identifies a strategy that maximizes the minimum possible payoff achievable against an optimal adversary.

6. (B) Alpha-beta cutoff.
Alpha-beta pruning is a fundamental optimization technique for the minimax search algorithm. It decreases the number of nodes evaluated in the search tree by identifying and discarding branches that cannot influence the final decision. By maintaining bounds on the possible values, the algorithm avoids exploring subtrees that are guaranteed to produce worse outcomes than previously found moves.

7. (A) P-3, Q-1, R-4, S-2
Various AI concepts serve distinct roles in problem-solving and intelligence testing. The Turing Test is essentially an imitation game for evaluating machine intelligence. Rational agents focus on maximizing performance measures. Heuristic search relies on evaluation functions to guide exploration efficiently. State spaces encompass every possible configuration or state that a problem can potentially occupy during search for a solution.

8. (A) reduce the number of game states to be evaluated in the search tree.
The primary benefit of employing alpha-beta pruning within a minimax search is the significant reduction in the total number of game states evaluated. By ignoring branches that are mathematically proven to be irrelevant to the final move selection, the search becomes much more efficient. This allows for deeper exploration of the game tree within the same time constraints.

9. (C) Scripts are primarily used for representing declarative, static facts.

NODIA

Scripts are specialized knowledge representation structures designed to handle stereotypical sequences of events occurring over time, such as dining at a restaurant. Unlike logic or semantic networks that often store static facts, scripts represent procedural and temporal knowledge. They provide context-specific expectations about the roles, props, and actions typically encountered during specific, frequently occurring social or physical interactions.

10. (A) Both A and R are true and R is the correct explanation of A.
In adversarial search, the alpha value represents the highest-value choice found so far by the maximizing player. During the search process at a MAX node, this value can only increase as better options are discovered. Because alpha acts as a lower bound for the maximizer's score, it ensures that the player is guaranteed at least that specific value.

11. (D) Sentence meaning using a small set of primitive action components.
Conceptual Dependency theory, developed by Roger Schank, aims to represent the underlying meaning of natural language sentences in a language-independent manner. It reduces complex linguistic expressions into a limited set of primitive actions and conceptual roles. This allows AI systems to perform inferences and answer questions based on the fundamental semantics of the input rather than its syntax.

12. (B) A formal, explicit specification of a shared conceptualization.
An ontology serves as a structural framework for representing knowledge within a specific domain. It provides a formal and explicit description of concepts, properties, and the relationships between them. By establishing a shared vocabulary, ontologies facilitate communication between different agents and systems, ensuring that information is interpreted consistently across various platforms and applications in the field of artificial intelligence.

13. (B) A, B, C, and D only
Knowledge representation involves several distinct formalisms. Semantic networks utilize graphs of nodes and arcs, while frames use slots and fillers for object attributes. Expert systems commonly employ IF-THEN rules for reasoning. Scripts are tailored for sequential events. However, ontologies are not restricted to procedural knowledge; they primarily focus on defining the fundamental concepts and relationships within a domain.

14. (D) Inference Engine
The inference engine is the core processing component of an expert system. It applies logical rules and reasoning strategies to the information contained within the knowledge base to derive new facts or conclusions. By matching current data against known rules, the engine simulates the human decision-making process, allowing the system to provide solutions or diagnoses for specific problems.

15. (A) The Sussman Anomaly, where a subgoal solution undoes a previous one.
Goal Stack Planning is a linear planning technique that can struggle when subgoals are interdependent. The Sussman Anomaly is a classic problem where achieving one goal requires undoing a previously satisfied goal, leading to inefficiency or failure. Non-linear planning avoids this by allowing actions for different subgoals to be interleaved, resolving conflicts through a more flexible ordering of operations.

16. (B) A representation of agent emotions and beliefs.
Classical planning systems focus on finding a sequence of actions to transition from an initial state to a goal state. These systems require formal representations of the world's states, available actions with their preconditions and effects, and the desired objectives. Subjective internal states such as emotions and personal beliefs are not typically included in these objective, goal-oriented computational models.

17. (B) Linear planning assumes subgoals are independent, while non-linear planning handles subgoal interactions.
Linear planning assumes that subgoals can be solved sequentially without interference, which often leads to failures like the Sussman Anomaly. Conversely, non-linear planning recognizes that actions taken to achieve one subgoal might affect others. By maintaining a partial ordering

Click Here to Buy book on Amazon

systems, a fact that is once proven remains true regardless of any subsequent axioms introduced to the knowledge base. This property ensures stability in reasoning but can be limiting when dealing with dynamic environments where information might change.

36. (D) Fuzzy sets allow for degrees of membership, represented by a value between 0 and 1.
Fuzzy set theory generalizes classical set theory by introducing the concept of partial membership. Unlike crisp sets, where an element either belongs or does not, fuzzy sets utilize a membership function that assigns a grade between zero and one. This allows for the representation of imprecise concepts, such as warm or tall, which do not have sharp boundaries in language.

37. (B) To model human-like reasoning in situations of uncertainty and imprecision.
The primary objective of fuzzy rule-based systems is to replicate human reasoning processes when dealing with vague or imprecise data. By utilizing linguistic variables and fuzzy logic, these systems can process information that is not easily quantified by traditional binary methods. This makes them highly effective for control systems and decision-making processes in complex, uncertain, or unpredictable real-world environments.

38. (C) A, B, C, and D only
Fuzzy logic is a multi-valued system designed to model the vagueness of human language. It employs fuzzifiers to convert crisp inputs into fuzzy sets and defuzzifiers to translate fuzzy outputs back into specific values. While it handles intermediate truth values effectively, it is actually highly suitable for controlling complex systems, contrary to the suggestion that it is less effective.

39. (D) Reproduction
Reproduction is the phase in a Genetic Algorithm where new candidate solutions are generated. This process involves selecting individuals from the current population based on their fitness and applying genetic operators to produce offspring. Through mechanisms like crossover and mutation, the algorithm creates new genetic variations, allowing the search to explore different regions of the potential solution space over generations.

40. (A) To evaluate how good a particular solution is in solving the problem.
The fitness function serves as the primary objective measure in a Genetic Algorithm. It quantifies the quality or performance of each individual chromosome within the population relative to the problem being solved. By assigning a numerical score to each candidate solution, the fitness function guides the selection process, ensuring that better solutions have a higher probability of reproducing.

41. (C) The representation of a potential solution as a chromosome (e.g., a bit string).
Encoding is a fundamental step in Genetic Algorithms that determines how a potential solution is represented as a chromosome. Common methods include binary encoding, where solutions are bit strings, or real-valued encoding. The choice of encoding significantly impacts the effectiveness of genetic operators like crossover and mutation, as it defines the mapping between the problem space and the genetic space.

42. (A) Crossover and Mutation
Crossover and mutation are the core operators in Genetic Algorithms responsible for balancing exploration and exploitation. Crossover combines successful traits from parent solutions to exploit known good regions of the search space. Mutation introduces random changes to ensure genetic diversity, allowing the algorithm to explore new areas and avoid getting trapped in local optima during the optimization process.

43. (B) P-4, Q-1, R-3, S-2
Genetic algorithms utilize specific operators to simulate natural evolution. Selection ensures that fitter individuals have a better chance of reproducing. Crossover merges genetic material from two parents to create offspring. Mutation introduces random variations to maintain diversity within the population. Encoding provides the specific format, such as bit strings or vectors, used to represent each candidate solution in the system.

44. (B) Supervised Learning
Supervised learning is a machine learning paradigm where a model is trained using a dataset containing labeled examples. Each input

NODIA

a hierarchical structure, typically a parse tree, according to the rules of a formal grammar. This process identifies how words relate to each other as subjects, verbs, and other grammatical constituents.

27. (D) Semantic Analysis
Semantic analysis is the level of Natural Language Processing that deals with the actual meaning of words and sentences. While syntax focuses on structural correctness, semantics attempts to map linguistic expressions to their underlying concepts and relationships in the real world. This process is essential for understanding the intent behind a sentence and enabling machines to process human information meaningfully.

28. (B) Pragmatics
Pragmatics is the branch of linguistics and Natural Language Processing that studies how context contributes to meaning. It goes beyond literal definitions to examine how social factors, speaker intent, and situational circumstances influence communication. Understanding pragmatics is crucial for AI systems to interpret nuances such as sarcasm, requests disguised as questions, and other subtle aspects of human-to-human language interaction.

29. (C) 2, 5, 1, 3, 4
Processing natural language typically follows a logical progression from small components to large contexts. The sequence begins with morphological analysis to identify word forms, followed by lexical analysis for tokenization. Next, syntactic analysis determines grammatical structure through parsing. Semantic analysis then interprets the literal meaning of the sentences. Finally, pragmatic analysis considers the broader context and intent behind the communication.

30. (D) Agents are proactive and autonomous, whereas objects only react to method calls.
While both agents and objects encapsulate state and behavior, their interaction models differ significantly. Objects are fundamentally reactive, executing code only when their methods are invoked by an external entity. In contrast, agents are autonomous and proactive, capable of taking initiative to achieve their own goals. Agents also possess social ability, allowing them to communicate and negotiate with other agents.

31. (A) The way the agent's percepts are mapped to its actions.
An agent's architecture defines its internal design and logic, specifically how it processes inputs from the environment and chooses appropriate responses. It represents the mapping from a sequence of percepts to the agent's actions. This structural framework determines whether the agent behaves as a simple reflex system, maintains an internal model, or acts based on complex goals and utilities.

32. (B) The communicative aspect.
The communicative aspect of multi-agent systems focuses on the interactions between individual entities. It encompasses the protocols, languages, and semantic frameworks that allow agents to exchange information, negotiate terms, and coordinate their actions. Effective communication is essential for agents to collaborate on shared tasks or resolve conflicts when operating within the same environment to achieve their respective goals.

33. (C) Agents work together, sharing information and tasks to achieve a common goal.
In a cooperative multi-agent system, individual agents are designed to work together harmoniously. They share relevant information and divide tasks to achieve a unified, common objective that benefits the entire group. This collaborative approach often involves negotiation and coordination to ensure that the combined efforts of all agents are more effective than the sum of their individual actions.

NODIA

34. (A) Competitive multi-agent system.
A scenario involving multiple ride-sharing cars from different companies vying for the same passengers exemplifies a competitive multi-agent system. In such environments, agents have conflicting goals or are competing for limited resources. Each agent acts to maximize its own utility, often at the expense of others, requiring strategies for negotiation or game-theoretic decision-making to succeed in the marketplace.

35. (A) Both A and R are true and R is the correct explanation of A.
Classic logic systems are generally monotonic, meaning the set of provable conclusions only grows as new information is added. In such

71. (D) The programming language of the sender agent.
In multi-agent communication, the fundamental components focus on the information exchange itself. This requires a sender, a receiver, and the actual content or message being transmitted. The specific programming language or internal implementation details of the agents involved are irrelevant to the communicative act, provided they adhere to the shared communication protocols and semantic standards of the multi-agent system.

72. (C) Sexual reproduction
In Genetic Algorithms, the crossover operator is directly inspired by sexual reproduction in biology. It involves taking genetic material from two parent solutions and combining them to create one or more offspring. By merging traits from successful parents, the algorithm hopes to produce a new generation of solutions that inherit and potentially improve upon the best characteristics of their genetic predecessors.

73. (B) Linear Planning (like early Goal Stack Planning)
The Sussman Anomaly is a famous problem in AI planning that highlights the deficiencies of strictly linear techniques like early Goal Stack Planning. In these systems, attempting to solve subgoals sequentially can lead to a situation where achieving one subgoal necessitates undoing another. This conflict demonstrates that many real-world problems require the interleaved action ordering found in non-linear planning methods.

74. (B) Conceptual Dependency : A framework for representing the meaning of sentences
Conceptual Dependency is a semantics-based framework designed to represent the underlying meaning of natural language sentences. It utilizes a set of primitive actions and roles to create a language-independent representation. While other concepts like alpha-beta cutoff or backpropagation are vital in AI, they relate to search optimization and neural network training, respectively, rather than unsupervised learning or general knowledge representation.

75. (A) A, B, and C only
A rational agent's decisions are based on its percepts, knowledge, and goals. Min-Max search is a recursive technique for navigating game trees. Ontologies establish the vocabulary and meaning for a specific domain. However, Goal Stack Planning is fundamentally a linear planning approach, not a form of hierarchical planning, which uses different levels of abstraction to simplify complex problem-solving tasks.

76. (C) To act as an auto-associative memory, retrieving a full pattern from a partial or noisy input.
The primary utility of a Hopfield Network lies in its ability to function as an auto-associative memory. It stores information as stable states within a recurrent network structure. When presented with incomplete or corrupted input data, the network's dynamics guide it toward the nearest stored pattern, effectively reconstructing the original information from the provided cues or noisy input signals.

77. (B) A function that estimates the cost from a node to the goal.
A heuristic function is a critical component of informed search algorithms. It provides an estimate of the remaining cost required to reach the goal state from a specific node in the search tree. By using problem-specific knowledge to guide the search toward more promising paths, heuristics significantly improve the efficiency of finding solutions compared to uninformed, blind search methods.

78. (A) parse tree or syntax tree.
The result of a parsing operation in Natural Language Processing is typically represented as a parse tree or a syntax tree. This hierarchical data structure illustrates the grammatical organization of a sentence, showing how words and phrases relate to each other according to the rules of a specific grammar. It identifies parts of speech and complex groupings like noun phrases.

79. (B) P-3, Q-1, R-2, S-4
AI techniques are suited for various problem types. Genetic algorithms are excellent for optimizing complex physical layouts like microchips. Expert systems excel at diagnosing specific failures based on known rules and symptoms. Fuzzy logic is ideal for managing imprecise control in home

Downloaded from www.pdf.tube

is paired with its corresponding correct output, allowing the network to learn the mapping between them. Through a process of error minimization, the model adjusts its internal parameters to accurately predict labels for new, unseen data in the future.

45. (C) Multilayer Perceptron (MLP)
The backpropagation algorithm is the standard method used for training multilayer perceptrons. It works by calculating the gradient of the loss function with respect to each weight in the network by applying the chain rule. By propagating the error backward from the output layer through the hidden layers, the algorithm allows the network to update its weights effectively and iteratively.

46. (D) Problems that are linearly separable.
A single-layer perceptron is a basic neural network model capable of solving classification problems that are linearly separable. This means the perceptron can only define a linear decision boundary, such as a line in two dimensions or a plane in three dimensions, to distinguish between different classes. It cannot solve non-linear problems like the XOR function without additional hidden layers.

47. (A) The goal is to discover hidden patterns or structures in unlabeled data.
Unsupervised learning is an approach where a neural network processes data without any pre-assigned labels or external teacher signal. The objective is to identify inherent structures, patterns, or clusters within the input space. Common applications include dimensionality reduction and data clustering, where the system organizes information based on similarities and differences discovered directly from the data points themselves.

48. (B) It acts as a content-addressable memory system.
A Hopfield network is a type of recurrent artificial neural network that functions as a content-addressable or associative memory system. It can store patterns and subsequently retrieve a full, stored pattern even when provided with a noisy or incomplete version of it. The network's state evolves toward a stable local minimum of an energy function, representing the stored information accurately.

49. (C) 2, 3, 1, 4
A typical Genetic Algorithm follows a specific cyclical process. It begins with the initialization of a population of candidate solutions. Next, the system performs an evaluation to calculate the fitness of each individual. This is followed by the selection of the best candidates for reproduction. Finally, crossover and mutation operators are applied to generate a new generation of candidate offspring.

50. (C) Reinforcement Learning
Reinforcement learning is a paradigm where an autonomous agent learns to make decisions by interacting with an environment. Unlike supervised learning, it does not rely on labeled data but instead uses a system of rewards and penalties. The agent's goal is to discover a policy that maximizes its cumulative reward over time through a process of trial and error interaction.

51. (B) A, C, and D only
Multilayer perceptrons are powerful models capable of solving complex, non-linear problems. Training these networks often involves the backpropagation algorithm, which calculates error gradients. A single perceptron typically employs a step function to produce binary outputs. However, self-organizing maps are fundamentally unsupervised learning models used for data visualization and clustering, rather than supervised learning as incorrectly suggested in some contexts.

NODIA

52. (B) A stereotyped sequence of events in a particular context.
In knowledge representation, a script is a structured framework designed to capture stereotypical sequences of actions and events within a specific situational context. For instance, a restaurant script defines the typical roles of customers and waiters, the expected order of events like ordering and paying, and the common props involved, allowing an AI to reason about various social interactions.

53. (B) Relationships between concepts.
Semantic networks represent knowledge through a graphical structure consisting of nodes

and directed arcs. The nodes typically represent specific objects, categories, or concepts within a domain. The arcs connect these nodes and represent the semantic relationships between them, such as is-a or has-part. This visual representation allows for efficient reasoning and property inheritance across the entire knowledge network.

54. (B) Acting humanly.
The Turing Test is the classic benchmark for the acting humanly approach to artificial intelligence. It focuses on the observable behavior of a machine rather than its internal thought processes or logical reasoning. If a machine can communicate in a way that is indistinguishable from a human to a judge, it is considered to have achieved human-level intelligence effectively.

55. (A) Both A and R are true and R is the correct explanation of A.
The STRIPS representation revolutionized AI planning by efficiently handling state transitions. By using add and delete lists, STRIPS only explicitly describes the facts that change when an action is executed. This implicitly assumes that everything not mentioned in those lists remains constant, effectively bypassing the complex frame problem that would otherwise require specifying every unchanged detail in the world.

56. (C) Propositional Logic
Propositional logic is a formal system that deals with absolute truth values, typically restricted to true or false. Because of this binary nature, it lacks the mechanisms to represent degrees of belief or probabilistic outcomes. Conversely, techniques like fuzzy logic, Bayesian probability, and certainty factors are specifically designed to manage the ambiguity and uncertainty inherent in real-world domain information.

57. (D) The specific value or data that is placed into a slot.
In frame-based systems, a frame represents a general concept or object, while slots represent its specific attributes or properties. A filler is the actual data or value assigned to a particular slot. For example, in a Car frame, Color might be a slot, and the filler would be a specific value like Red for a particular instance.

58. (A) Linear planners produce a strictly ordered sequence of actions, while non-linear planners produce a partially ordered set.
The fundamental distinction between linear and non-linear planning lies in action ordering. Linear planners generate a total order, where actions are executed in a strict, single sequence. Non-linear planners, however, create a partial ordering, allowing actions to be interleaved. This flexibility enables non-linear systems to solve complex problems with interacting subgoals that often defeat strictly sequential linear planning systems.

59. (A) P-3, Q-4, R-1, S-2
Natural language tasks cover different aspects of language understanding. Semantic analysis derives meaning independent of context, while pragmatic analysis interprets intended meaning based on the specific situation. Syntactic analysis focuses on grammatical structure using formal rules. Morphological analysis examines the internal structure of words, identifying roles such as tense or plurality based on word forms and their constituent parts.

60. (B) The success of one agent often comes at the expense of another.
In a competitive multi-agent system, the individual goals of the participating agents are often in direct conflict. This is typical of zero-sum environments where resources are limited and one agent's gain results in another's loss. Agents in these systems must employ strategic decision-making, such as game-theoretic models, to maximize their own utility while anticipating moves of adversaries.

61. (C) The support of a fuzzy set includes elements with a membership grade of exactly 0.
Fuzzy sets are defined by membership functions that assign grades between zero and one. Standard operations include union via maximum and intersection via minimum values. The complement is calculated by subtracting the membership from one. Crucially, the support of a fuzzy set is the set of elements whose membership grade is strictly greater than zero, excluding those with no membership value.

62. (B) A, B, C, and E only

Click Here to Buy book on Amazon

Genetic Algorithms are stochastic optimization techniques inspired by biological evolution, utilizing a population of candidate solutions. They rely on probabilistic transition rules and primary operators like crossover and mutation. Unlike some gradient-based methods, they do not require the search space to be differentiable or continuous, making them highly versatile for solving discrete and complex optimization problems in various domains.

63. (B) propagating the error backwards through the network to update the weights.
Backpropagation is a vital learning algorithm for training multi-layer artificial neural networks. After a forward pass generates an output, the algorithm calculates the difference between the prediction and the target. This error is then propagated backward through the layers using the chain rule to determine how much each weight contributed to the error, allowing for precise adjustments of the weights.

64. (C) dimensionality reduction and visualization of high-dimensional data.
A Self-Organizing Map is a type of artificial neural network trained using unsupervised learning. Its primary function is to create a low-dimensional, typically two-dimensional, representation of complex, high-dimensional input data. By preserving the topological properties of the original space, SOMs are highly effective for data visualization, clustering, and identifying relationships within large, multifaceted datasets across various computational research fields.

65. (A) 1, 2, 3, 4
The STRIPS planning algorithm follows a recursive decomposition process. It starts by pushing the main goal onto the stack. If the top item is a compound goal, it is broken down into simpler subgoals. If an action is at the top, it is popped and replaced by its preconditions. Once a single goal matches the current state, it is satisfied.

66. (B) The realization that a branch of the game tree can be pruned if a move is found that is worse than a previously examined move.
The alpha-beta procedure is based on logical pruning of the search tree. If the algorithm discovers a move that is already worse for a player than a move found earlier, it can safely stop evaluating the current branch. This is because rational players will not choose that path, making further exploration unnecessary and significantly increasing search efficiency in competitive games.

67. (D) Script
A script is the most appropriate knowledge representation for standard, stereotypical sequences of events like visiting a restaurant. It provides a structured framework that includes expected roles, props, and a timeline of actions. This allows an AI to understand and predict common social interactions and fill in missing information based on the typical story associated with the specific situational context.

68. (A) Maximizing a performance measure given the available information.
The rational agent approach defines intelligence through the lens of objective performance. Instead of focusing on human-like imitation or internal thinking processes, it evaluates an agent based on its actions. A rational agent acts to achieve the best possible or expected outcome by maximizing its performance measure, considering the information it has perceived from the environment and its built-in knowledge.

69. (B) Fuzzification
Fuzzification is the initial step in a fuzzy inference system. It involves taking a crisp, numerical input value and determining its degree of membership in various fuzzy sets. For instance, a temperature of 22 degrees Celsius might be mapped to a membership of 0.7 in the warm set and 0.3 in the cool set, allowing the system to process imprecise data.

70. (A) Both A and R are true and R is the correct explanation of A.
The universal approximation theorem states that a multilayer perceptron can approximate any continuous function given sufficient hidden neurons. This capability is fundamentally derived from the inclusion of non-linear activation functions in the hidden layers. Without these non-linearities, the network would essentially be a linear model, incapable of capturing the complex, curved relationships present in many real-world datasets and functions.

NODIA

of actions and allowing interleaving, non-linear planners can resolve conflicts and dependencies more effectively than their strictly linear counterparts.

18. (A) A stack
Goal Stack Planning utilizes a stack data structure to manage its objectives and sub-objectives. The system pushes the primary goal onto the stack and subsequently decomposes it into necessary preconditions or subgoals. This Last-In, First-Out mechanism allows the planner to systematically address the most immediate requirements first, though it may struggle with complex interactions between different goals on the stack.

19. (A) Both A and R are true and R is the correct explanation of A.
Hierarchical planning is a powerful technique that addresses complex problems by decomposing them into multiple levels of abstraction. By first developing a high-level plan that ignores minor details, the system manages complexity more effectively. Once the abstract structure is established, the planner refines each step into more specific actions, ensuring that the final detailed plan is logically consistent.

20. (D) Preconditions, Add List, and Delete List.
In the STRIPS planning formalism, an action is defined by three essential lists. Preconditions specify the facts that must be true for the action to be performed. The add list contains the new facts that become true after the action occurs. The delete list identifies facts that are no longer true, allowing the system to track state changes effectively and efficiently.

21. (B) POP commits to a total ordering of all actions from the very beginning.
Partial Order Planning is a non-linear approach that avoids making unnecessary ordering decisions. Instead of committing to a strict sequence from the start, it maintains only the constraints essential for resolving conflicts or meeting preconditions. This flexibility allows the planner to explore different action sequences and only establishes a total order when the plan is finalized and ready for execution.

22. (A) P-4, Q-3, R-1, S-2
Planning techniques in AI are categorized by their underlying logic and structures. Goal Stack Planning relies on a LIFO stack to manage subgoals. Hierarchical planning uses different levels of abstraction to simplify complex tasks. The STRIPS formalism represents actions through specific precondition and effect lists. Partial Order Planning operates by searching through the space of possible plans rather than states.

23. (D) It provides an objective and mathematically analyzable criterion for success.
The rational agent approach defines intelligence as the ability to select actions that maximize expected performance. Unlike the Turing Test, which relies on subjective human judgment and imitation of behavior, the rational agent model provides a clear, objective framework. This allows for mathematical analysis and formal verification of an agent's performance, making it more robust for diverse engineering applications.

24. (A) return an optimal solution if one exists.
In search algorithms, admissibility is a crucial property related to heuristic functions. An admissible heuristic never overestimates the actual cost to reach the goal. When used in algorithms like A*, an admissible heuristic ensures that the first solution found is the optimal one. This guarantee makes it a highly desirable characteristic for finding the most efficient paths in problem-solving.

25. (C) A, B, C, and D
Min-Max search systematically explores game trees using depth-first techniques. Alpha-beta pruning optimizes this by eliminating unnecessary branches, with its efficiency heavily dependent on move ordering. Alpha serves as the lower bound for the maximizing player, while beta acts as the upper bound for the minimizing player. These components together form the basis for efficient adversarial search in game-playing artificial intelligence.

26. (A) Analyzing a string of symbols to determine its grammatical structure.
Parsing is a vital stage in Natural Language Processing that focuses on the structural organization of language. It involves taking a sequence of tokens and arranging them into

Downloaded from www.pdf.tube

appliances. NLP parsers are specifically designed to analyze the grammatical correctness and structure of natural language sentences correctly.

80. (B) The autonomous and goal-directed nature of the components (agents).
While distributed systems involve multiple processors working together, multi-agent systems are distinguished by the autonomy of their components. Each agent in the system is a goal-directed entity capable of making independent decisions to achieve its specific objectives. Agents can react to changes, take initiative, and interact socially, which goes beyond simple communication and shared processing found in distributed architectures.

81. (C) To maintain genetic diversity in the population and prevent premature convergence.
In Genetic Algorithms, the mutation operator introduces small, random changes to individual chromosomes. Its primary role is to ensure genetic diversity within the population. By occasionally creating entirely new traits that were not present in the initial population, mutation allows the algorithm to explore new regions of the search space and prevents it from settling too quickly on suboptimal solutions.

82. (A) Always true.
A multilayer perceptron derives its power from the interaction between multiple layers of neurons. If the hidden layers are removed, the architecture simplifies to a direct mapping from input to output. This configuration is functionally identical to a single-layer perceptron, which can only solve linearly separable problems. Regardless of the activation function used, the underlying mathematical model remains linear.

83. (C) Self-organizing Maps (SOM)
Self-organizing maps are a prominent example of unsupervised learning in neural networks. Unlike supervised methods like backpropagation or the perceptron learning rule, which require labeled data, or reinforcement learning, which uses rewards, SOMs discover the underlying structure of the data on their own. They create a low-dimensional topological representation of input data, making them ideal for clustering and visualization tasks.

84. (D) It always produces a plan that is a strictly linear sequence of actions.
Partial Order Planning is characterized by its flexibility in action sequencing. It starts with minimal constraints and adds ordering only to resolve specific threats or meet preconditions. Unlike linear planners, POP does not produce a single sequence but rather a partially ordered set of actions. This allows for parallel execution or flexible scheduling, only finalizing the total order when necessary.

85. (A) Both A and R are true and R is the correct explanation of A.
Fuzzy logic is an essential tool for handling the imprecision found in human language and real-world data. Unlike classical logic, which requires strict binary belonging, fuzzy logic allows for graded membership. This enables systems to represent and reason with concepts like tall or hot more naturally, as these terms do not have sharp, well-defined boundaries in common human usage.

86. (A) State space defines the structure of a problem for search, while knowledge representation is about storing and reasoning with facts about the world.
State space representation is a problem-solving framework that defines the possible states and transitions for a specific search task. Knowledge representation is a broader field focused on how to formally represent, store, and manipulate facts and relationships about the world. While state spaces are used for planning and search, knowledge representation provides the underlying semantics and basis for intelligent reasoning.

87. (D) Both (A) and (B) are required.
The Min-Max algorithm guarantees finding an optimal move under specific conditions. First, the game must be a two-player, finite, zero-sum game with opposing goals. Second, the algorithm must explore the entire game tree down to the terminal nodes to calculate exact utility values. Without complete exploration or if the game structure is different, the result is an estimate.

88. (C) A, B, C, and D only
STRIPS provides a formal language for planning, while Conceptual Dependency offers primitives

NODIA

for semantic representation. Hopfield networks are specialized recurrent neural models, and rational agents are defined by their goal-oriented actions. However, the Turing Test is not considered a definitive measure of consciousness; it evaluates observable behavior and linguistic indistinguishability, leaving the question of internal subjective experience entirely unanswered.

89. (A) Top-down parsing
Top-down parsing is a technique that constructs a parse tree by starting from the most abstract symbol, the start symbol of the grammar. The algorithm then systematically expands the non-terminal symbols according to grammar rules until it reaches the leaves, which must match the actual words in the input sentence. This approach predicts the structure of the sentence before verification.

90. (A) Ensuring effective coordination and communication to achieve the common goal.
In cooperative multi-agent systems, the primary difficulty lies in synchronization. Agents must effectively communicate their states, intentions, and discovered information to others. Ensuring that the collective behavior of autonomous entities remains coordinated toward a shared objective requires robust protocols and strategies for task allocation, preventing redundant efforts and managing conflicts that might arise during the collaborative problem-solving process.

91. (A) 2, 1, 3
The structural complexity of neural networks progresses from simple single-layer units to sophisticated architectures. A perceptron is the most basic form, consisting of a single neuron. The multilayer perceptron introduces hidden layers, allowing for more complex mappings. Finally, deep neural networks utilize many hidden layers, significantly increasing the model's capacity to learn hierarchical features and intricate patterns from large datasets.

92. (B) The facts that become true in the world after the operator is applied.
In the STRIPS planning formalism, the add list is a crucial component of an action's definition. It specifies the set of propositions or facts that transition from false to true in the environment immediately after the action is executed. This list, along with the delete list, allows the planner to update its internal model of the world's state following each step.

93. (B) the design that maps an agent's percepts to its actions.
Agent architecture refers to the fundamental design of an intelligent system. It describes the structural framework and logic that dictates how the agent processes information perceived from its environment and decides on subsequent actions. This architecture serves as the link between the agent's sensors and actuators, determining its behavior and ability to achieve goals within its specific operational and situational environment.

94. (D) Converting the final fuzzy output set into a single crisp value.
In the context of fuzzy logic control, defuzzification is the critical final stage. After applying fuzzy rules to produce a combined fuzzy output, the system must generate a specific, actionable command. Defuzzification translates the theoretical fuzzy result back into a single, crisp numerical value that can be directly used to control an actuator or make a final decision in the environment.

95. (A) Both Statement I and Statement II are true.
The fitness function is the driving force behind the evolutionary process in a Genetic Algorithm. It must accurately represent the quality of potential solutions relative to the problem's goals. If the function is poorly defined or fails to capture the essential objectives, the algorithm may optimize for irrelevant traits, leading to suboptimal or entirely nonsensical results that do not solve the actual problem.

96. (A) plans at different levels of abstraction, refining high-level plans into more detailed ones.
Hierarchical planning is a strategy used to manage complex tasks by decomposing them into multiple levels. The planner first creates a high-level, abstract plan that captures the broad steps needed to reach the goal. Each abstract step is then refined into progressively more detailed sub-actions. This hierarchical approach significantly reduces the search space and makes large-scale problem-solving more manageable for agents.

Click Here to Buy book on Amazon

97. (B) the meaning of sentences.
Semantic analysis is the stage of Natural Language Processing focused on identifying the literal meaning of linguistic expressions. While syntax deals with structural rules and pragmatics focuses on context, semantics aims to understand the relationships between words and the concepts they represent. This process is essential for extracting information from text and enabling machines to perform reasoning based on human language input.

98. (D) An agent cannot encapsulate state, only behavior.
Both agents and software objects are designed to encapsulate both state and behavior. The primary difference lies in their interaction model and autonomy. While objects react to external stimuli through method calls, agents are proactive and autonomous entities. Suggesting that an agent lacks state encapsulation is incorrect, as agents must maintain internal information to track their environment and achieve their specific goals.

99. (A) P-4, Q-2, R-3, S-1
Knowledge representation methods use distinct structural elements. Logic systems are built on predicates and terms. Semantic networks utilize nodes and arcs to represent concepts and relationships. Scripts organize information into temporal scenes and events. Frame-based systems employ a slot-and-filler structure, where slots define attributes of a concept and fillers provide the specific values or pointers to other related frames.

100. (B) They allow for graded or partial membership.
Fuzzy set theory is uniquely capable of representing concepts with imprecise boundaries. By assigning membership values between zero and one, fuzzy sets allow for the idea of graded belonging. This reflects how humans naturally categorize things like height; a person isn't just tall or not tall, but can be somewhat tall or very tall, providing a more nuanced and realistic model.

101. (B) A, B, and C only
Supervised learning relies on labeled data, while unsupervised learning discovers patterns in unlabeled datasets. Reinforcement learning involves an agent learning from rewards and penalties through environmental interaction. However, the backpropagation algorithm is a technique used specifically for supervised learning to minimize error in multi-layer networks, rather than being an example of reinforcement learning as incorrectly suggested in some specific contexts.

102. (A) the best possible score (lowest value) for the minimizing player found so far in the search.
In the alpha-beta pruning algorithm, the beta value serves as an upper bound on the possible score for a node. It represents the best, or lowest, value found so far by the minimizing player along the current search path. When a value is found that exceeds beta, the branch can be pruned because the minimizing player will never choose that path.

103. (C) Rules
Rule-based representation is a common formalism in AI that uses if-then structures to represent knowledge and logical deductions. This approach is the foundation for most classical expert systems, where a set of rules is applied to a collection of facts by an inference engine to derive new information or reach a specific conclusion, mimicking the reasoning of human specialists accurately.

104. (A) Both A and R are true and R is the correct explanation of A.
A single-layer perceptron is limited to solving problems that are linearly separable. The XOR function is a classic example of a non-linearly separable problem, as its outputs cannot be divided by a single straight line or hyperplane. Consequently, a single perceptron cannot learn this function, necessitating the use of multilayer networks with hidden layers to achieve non-linear decision boundaries.

105. (B) an action added to the plan could potentially negate a required condition (causal link) of another action.
In Partial Order Planning, a threat occurs when a new action's effects conflict with a causal link established between two other actions. If an action negates a precondition that another action depends on, it threatens the validity of the plan. The planner must resolve these threats by adding ordering constraints, such as ensuring

Downloaded from www.pdf.tube

140. (C) XOR gate
The XOR problem is a famous example of a task that is not linearly separable, meaning its outputs cannot be divided by a single straight line. While simple AND or OR gates can be learned by a single-layer perceptron, the XOR function requires a multilayer perceptron with at least one hidden layer to create the necessary non-linear decision boundary.

141. (C) A, B, C, and D only
Planning in AI involves several methodologies. Non-linear planning manages interleaved subgoals, while Goal Stack Planning uses a LIFO strategy. Hierarchical planning refines abstract goals, and STRIPS uses specific precondition-add-delete lists. While these are all valid techniques, not all planning algorithms are guaranteed to find an optimal plan; optimality depends on the specific search strategy and heuristics employed.

142. (B) exhibit human-like intelligence in conversation.
The Turing Test focuses on a machine's ability to demonstrate intelligent behavior through natural language interaction. The test is passed if a machine can converse so convincingly that a human judge cannot distinguish it from another human. This benchmark prioritizes indistinguishable behavioral imitation over internal logic or raw computational power as a measure of a machine's overall intelligence.

143. (A) the guaranteed minimum score for the MAX player on that path so far.
In the alpha-beta search algorithm, the alpha value represents a lower bound on the score for the maximizing player. As the algorithm explores a MAX node, it updates alpha whenever it finds a move that leads to a higher value. This value represents the best score the MAX player is already assured of achieving along that specific branch of search.

144. (A) a formal specification of the concepts and relationships in a domain.
An ontology is a structural knowledge representation that defines the fundamental concepts and relationships within a specific area of interest. It provides a formal, explicit specification of a shared conceptualization, establishing a common vocabulary for agents and systems. This standardization ensures that information is understood consistently across different platforms, facilitating interoperability and automated reasoning in various domains.

145. (A) Both Statement I and Statement II are true, and Statement II is the correct explanation of Statement I.
Partial Order Planning is often more efficient for complex tasks because it follows the principle of least commitment. Instead of deciding on a total order for all actions immediately, it only imposes constraints when necessary to resolve threats or fulfill preconditions. This flexibility prevents the planner from making arbitrary ordering choices that might later lead to conflicts, saving significant computational effort.

146. (B) a universal Turing machine for general computation.
Typical agent architectures are designed around environmental interaction, consisting of sensors for perception, actuators for action, and an agent program to map percepts to responses. While an agent may utilize computational processes, a universal Turing machine is an abstract mathematical model of computation and is not a standard component or requirement for defining the architecture of an intelligent agent.

147. (A) agents work together towards a common objective.
In a cooperative multi-agent system, individual autonomous entities coordinate their actions and share information to achieve a shared, overarching goal. These agents are designed with compatible objectives, and the success of the system depends on their ability to collaborate effectively. This teamwork allows the group to solve complex problems that would be beyond the capabilities of any single agent.

148. (C) To combine genetic material from two parent chromosomes to create offspring.
The crossover operator is a fundamental mechanism in Genetic Algorithms used during the reproduction process. It takes two high-quality parent solutions and exchanges sections of their genetic code to produce new offspring. By combining the successful traits of both parents, crossover enables the algorithm to exploit known

NODIA

Genetic Algorithms are heuristic search and optimization techniques inspired by the biological process of evolution. They operate on a population of candidate solutions, applying operators like selection, crossover, and mutation over multiple generations. By favoring the most successful individuals for reproduction, the algorithm simulates natural selection, gradually improving the population's overall fitness and converging toward an optimal or near-optimal solution.

115. (B) B, C, D, and E only
Modern AI concepts cover a wide range of functions. Ontologies define domain vocabularies, and Min-Max search provides a depth-first strategy for games. Fuzzy sets enable graded membership, while backpropagation facilitates supervised learning in neural networks. However, the Turing Test is specifically designed to evaluate a machine's ability to act humanly through conversation, rather than its ability to act rationally in environments.

116. (B) The machinery that the agent executes on, providing percepts and receiving actions.
An agent's behavior is the result of the interaction between its program and its architecture. The architecture represents the physical or virtual machinery that provides the agent with its perceptions from the environment and executes the chosen actions. It serves as the bridge between the high-level decision-making logic of the agent program and the actual environment in which the agent operates.

117. (C) conflicting, such that one agent's gain is another's loss.
In competitive multi-agent systems, agents are often in a zero-sum or similar relationship. Their goals are in direct conflict, meaning that the achievement of one agent's objective typically hinders or prevents another's success. This environment requires agents to use strategic reasoning, such as game theory, to anticipate opponent actions and maximize their own utility in the face of environmental adversity.

118. (C) It determines how a potential solution is represented (e.g., as a binary string or a real-valued vector).
Encoding is a foundational step in Genetic Algorithms that defines the mapping between a problem solution and its representation as a chromosome. Whether using binary strings, permutations, or real-valued vectors, the chosen encoding method dictates how the genetic operators like crossover and mutation will interact with the solution space. A well-designed encoding is essential for the algorithm to search effectively.

119. (A) one or more hidden layers between the input and output layers.
The structural difference between a simple perceptron and a multilayer perceptron lies in the hidden layers. While a single perceptron passes inputs directly to an output node, an MLP includes intermediate layers of neurons. These hidden layers, combined with non-linear activation functions, allow the MLP to model complex, non-linear relationships that a basic, single-layer perceptron simply cannot represent effectively.

120. (B) Unsupervised learning
Self-Organizing Maps are a type of neural network trained through unsupervised learning. This means they process input data without any pre-defined labels or target outputs. SOMs work by adjusting their internal weights to discover the underlying structure and topology of the input space, creating a simplified, low-dimensional map that clusters similar data points together for visualization and unsupervised analysis.

121. (A) It selects actions to maximize its expected performance measure.
A rational agent is characterized by its goal-oriented behavior. It does not necessarily mimic human decision-making or possess perfect knowledge; instead, it uses the information available to it through its percept sequence to select actions that are most likely to achieve the best possible outcome. This maximization of the expected performance measure is the core definition of rationality in AI.

122. (A) Breadth-First Search
Breadth-First Search is an example of an uninformed or blind search technique. It explores the search space systematically by visiting all nodes at a given depth before moving to the next

Click Here to Buy book on Amazon

level. Unlike heuristic search methods such as A* or Greedy Best-First Search, BFS does not use any problem-specific information or evaluation functions to guide its exploration toward the goal.

123. (C) the facts from the current state that are no longer true after the operator is applied.
The delete list is a vital part of an action's definition in the STRIPS planning model. It specifies the set of conditions or facts that cease to be true once the action is completed. By removing these items from the current state representation while adding those from the add list, the planner can accurately track how each action modifies the world state.

124. (B) A knowledge base and an inference engine.
A classic expert system is composed of two primary functional units. The knowledge base contains the specialized domain information, typically represented as a set of facts and rules. The inference engine is the reasoning component that applies these rules to current data to reach conclusions. Together, they allow the system to simulate the specialized decision-making processes of a human domain expert.

125. (D) A, B, C, and D
Knowledge representation utilizes various structured formats. Semantic networks use graphs, frames organize stereotypical situations into slots and fillers, and scripts model sequences of events. Additionally, Conceptual Dependency theory provides a system of primitive actions for semantic modeling. Each of these techniques is a valid and correct method for representing information in a way that an AI system can use for reasoning.

126. (B) adjust the weights in a neural network to minimize error.
The backpropagation algorithm is the primary method for training multi-layer neural networks. It operates by calculating the gradient of the loss or error function with respect to each weight in the network. By propagating the prediction error backward from the output layer, the algorithm determines how to adjust weights to reduce the difference between the network's output and target value.

127. (B) To estimate the expected utility of the game from that state.
In game tree search, exploring every possible move to the end of the game is often computationally impossible. When a search reaches a predefined cutoff depth, it uses an evaluation function to provide a heuristic estimate of the utility of that state. This function approximates how favorable the current board configuration is for a player without continuing to the terminal nodes.

NODIA

128. (B) To delay making ordering decisions until forced to, maintaining flexibility.
Partial Order Planning is based on the principle of least commitment. It avoids establishing a strict, total sequence of actions from the beginning. Instead, it only imposes ordering constraints when they are necessary to resolve conflicts or ensure preconditions are met. This flexibility allows the planner to efficiently manage complex tasks with independent subgoals and find solutions without premature ordering errors.

129. (A) Both A and R are true and R is the correct explanation of A.
Effective interaction in multi-agent systems requires a shared communication framework. Agents must use a common language and adhere to specific protocols so that messages are understood consistently across the system. Without this shared understanding of the syntax and semantics of communication, autonomous agents would be unable to share information, negotiate terms, or coordinate their actions to achieve collective goals.

130. (C) It uses a membership function that assigns a degree of membership between 0 and 1 to elements.
The defining feature of a fuzzy set is its ability to represent partial membership. Unlike a classical crisp set where an element either belongs or does not, a fuzzy set uses a membership function to assign a value between zero and one. This mathematical model allows for the representation of imprecise concepts, reflecting the gradual transitions found in many real-world categories.

131. (A) Selection, Crossover, and Mutation
In Genetic Algorithms, the reproduction phase creates new candidate solutions using three core operators. Selection identifies the fittest

individuals to act as parents. Crossover then combines the genetic material of these parents to create offspring. Finally, mutation introduces random changes to the offspring's genetic code. This cycle ensures that the population evolves over time toward increasingly optimal solutions to the problem.

132. (A) Syntactic Analysis (Parsing)
Syntactic analysis, commonly referred to as parsing, is the stage of Natural Language Processing that identifies the grammatical structure of a sentence. It involves breaking the text into constituent parts like noun and verb phrases based on formal grammar rules. This process determines the logical relationships between words, which is a necessary precursor to understanding the sentence's actual meaning or intent.

133. (D) A fuzzy inference engine
Classical planning systems are formally defined by a set of possible world states, a set of actions with clear preconditions and effects, and a specific goal condition. They focus on finding discrete sequences of operations. While fuzzy inference engines are important in AI for handling imprecision and control, they are not a standard component of the classical planning framework itself.

134. (B) The protocols and languages agents use to interact.
The communicative aspect of a multi-agent system encompasses the social interactions between individual autonomous entities. It defines the formal languages, semantic frameworks, and interaction protocols that allow agents to exchange information and intent. This structured communication is the foundation for all collaborative or competitive behaviors, enabling agents to negotiate over resources and coordinate their actions effectively in environments.

135. (B) A set of fuzzy sets with corresponding degrees of membership.
Fuzzification is the translation of precise data into the fuzzy logic domain. The output of this stage is a collection of membership values that describe how much the input belongs to various overlapping fuzzy sets. For instance, a single speed input might be fuzzified into degrees of membership in sets like slow, medium, and fast, which are then processed by rules.

136. (A) candidate solutions to a problem.
A Genetic Algorithm operates on a population of individuals, where each individual represents a potential solution to the problem being solved. Through generations of selection and genetic variation, the algorithm evolves these candidate solutions to improve their fitness. This population-based approach allows the search to explore multiple areas of the solution space simultaneously, increasing the likelihood of finding a global optimum.

137. (C) Hopfield Network
While some neural networks like Multilayer Perceptrons require labeled data for training via backpropagation, others operate without external guidance. A Hopfield Network is an example of a recurrent network that functions as an associative memory through unsupervised processes. It learns to store patterns by adjusting weights based on the patterns themselves, allowing it to later reconstruct full information from partial inputs.

138. (B) Supervised learning has labeled data with correct outputs, while reinforcement learning uses a reward signal.
The fundamental difference between supervised and reinforcement learning lies in the nature of the feedback provided during training. Supervised learning uses a teacher signal in the form of labeled examples with known correct outputs. In contrast, reinforcement learning involves an agent interacting with an environment and receiving a scalar reward or penalty signal based on the quality of its chosen actions.

139. (A) P-4, Q-3, R-2, S-1
Genetic algorithm terminology is inspired by biology. A chromosome is the encoded representation of a single solution, while a gene is a part of that chromosome representing a specific parameter. The population is the entire set of solutions currently being evolved. Fitness is the essential measure of how well a particular solution satisfies the objectives of the problem being solved.

Downloaded from www.pdf.tube

cannot make decisions in such scenarios is incorrect, as that is their purpose.

158. (A) P-3, Q-4, R-1, S-2
Different neural architectures serve specific functions. Perceptrons are simple linear classifiers. Multilayer perceptrons with hidden layers act as universal function approximators. Hopfield networks are designed for content-addressable memory through recurrent connections. Self-organizing maps are used for unsupervised clustering and creating low-dimensional representations of data. Each of these models provides unique capabilities for solving diverse types of learning problems.

159. (C) randomly exploring new parts of the search space.
In a Genetic Algorithm, the mutation operator introduces random variations to individual chromosomes. This function is most similar to the process of exploration in search. By creating traits that may not have been present in the initial population, mutation ensures that the algorithm can discover new areas of the search space, preventing it from getting trapped in suboptimal local solutions.

160. (B) Agents are a specific type of object with added autonomy.
While agents and objects share characteristics like encapsulation of state and behavior, agents possess additional capabilities. Agents are autonomous entities that can take initiative to achieve their goals, whereas objects are fundamentally reactive and only execute actions when their methods are called. In this sense, agents can be viewed as specialized, more sophisticated versions of software objects that exhibit proactive behavior.

161. (B) A and C only
Alpha-beta pruning is a logical optimization of the minimax algorithm that eliminates irrelevant branches of the search tree. Crucially, it never changes the final move chosen, as it only removes subtrees that cannot influence the decision at the root. The efficiency of this technique is highly dependent on move ordering, with better orderings leading to more significant pruning.

162. (A) To apply the fuzzy rules to the fuzzy inputs to produce a fuzzy output.
The inference engine is the decision-making heart of a fuzzy logic system. Its role is to take the fuzzy sets generated from inputs and apply the collection of fuzzy if-then rules to them. This process determines the degree to which each rule is satisfied and combines the results to produce a fuzzy output, simulating the human reasoning process accurately.

163. (B) The maintenance of a population of candidate solutions.
A central component of the Genetic Algorithm's working principle is the use of a population of candidate solutions. Unlike search algorithms that follow a single path, a GA maintains multiple potential solutions simultaneously. This diversity allows the algorithm to explore different regions of the search space in parallel, reducing the risk of getting stuck in a single local optimum.

164. (C) The gradient of the error function with respect to the network's weights.
The backpropagation algorithm is a technique for supervised learning in neural networks. Its primary function is to calculate the gradient of the loss or error function for each weight in the network. By applying the chain rule of calculus from the output layer backward, it identifies how each individual weight contributes to the overall prediction error during training.

165. (C) auto-association and optimization problems.
Hopfield networks are a specialized type of recurrent artificial neural network. They are primarily utilized for auto-association tasks, where they can store patterns and retrieve them from noisy or incomplete inputs. Additionally, their energy-minimization dynamics make them useful for solving complex optimization problems, such as the traveling salesman problem, by mapping the problem constraints to the network's energy function.

166. (C) an internal state that represents a model of the world.
The key distinction of a model-based reflex agent is its ability to maintain an internal state. This state acts as a model of the world, allowing the agent to track parts of the environment that are not currently visible to its sensors. This internal

NODIA

good areas of the search space, potentially creating superior solutions in subsequent generations.

149. (A) 1, 3, 2
The development of neural networks followed a distinct historical path. The perceptron, introduced in the late 1950s, was the first simple model. In the early 1980s, John Hopfield introduced his recurrent network for associative memory. Subsequently, the mid-1980s saw the widespread popularization of the backpropagation algorithm, which finally allowed for the efficient training of multilayer networks with hidden layers.

150. (B) a fuzzy inference rule.
A rule that links linguistic variables using logical operators, such as IF temperature is HIGH AND humidity is HIGH, THEN fan_speed is FAST, is a classic example of a fuzzy inference rule. These rules form the core of a fuzzy rule-based system, allowing the AI to perform human-like reasoning by mapping imprecise fuzzy inputs to corresponding fuzzy outputs.

151. (D) The thinking rationally approach (Laws of Thought).
The thinking rationally approach in AI is rooted in the Laws of Thought tradition. It focuses on using formal logic and syllogisms to build systems that reason according to rigorous, provable rules. This approach emphasizes the correctness of the internal thought process, ensuring that the machine arrives at logical conclusions based on a formal representation of the world's knowledge.

152. (B) h(n) never overestimates the true cost to the goal.
In search theory, a heuristic function is defined as admissible if it provides an optimistic estimate of the remaining cost to reach the goal. This means the value it returns is always less than or equal to the actual minimum cost. Admissibility is a critical property because it ensures that algorithms like A* will always find the optimal solution.

153. (D) A specific attribute or property of the concept being represented.
Within the frame-based knowledge representation system, a frame represents a general concept or object. A slot is an individual component of that frame that specifies a particular attribute, property, or characteristic of the concept. Each slot can then be assigned a filler, which is the specific value or data that describes that attribute for a particular instance.

154. (A) Both A and R are true and R is the correct explanation of A.
Scripts are highly effective for natural language understanding in stereotyped contexts because they provide a pre-defined framework of expectations. By organizing information into standard roles, props, and sequences of events, scripts allow an AI to infer details that are not explicitly stated in a narrative. This ability to fill in gaps based on typical situations is essential for comprehension.

155. (C) It reduces complexity by decomposing the problem into hierarchical levels.
The primary advantage of hierarchical planning is its ability to handle large-scale, complex problems by breaking them into manageable levels of abstraction. By first establishing a high-level plan and then refining each abstract step into detailed sub-actions, the planner avoids searching through an overwhelming number of low-level states. This structural decomposition significantly increases the efficiency of the planning process.

156. (D) Lexical Analysis
Lexical analysis is the foundational stage of Natural Language Processing. It involves taking a raw string of text and identifying its basic constituent parts, such as individual words, numbers, and punctuation marks. This process, also known as tokenization, is essential because it prepares the text for more complex linguistic operations like identifying grammatical structure or deriving actual semantic meaning later.

157. (D) A utility-based agent cannot make decisions when goals are conflicting.
Intelligent agents are categorized by their decision-making logic. Simple reflex agents respond to current percepts, while model-based ones track the world. Goal-based agents act to reach specific states. However, utility-based agents are specifically designed to handle conflicting goals by using a utility function to weight different outcomes. Suggesting they

Click Here to Buy book on Amazon

the threatening action occurs before or after the linked actions.

106. (B) coordination, cooperation, and negotiation among agents.
The communicative aspect in a multi-agent system is vital for enabling interaction between autonomous entities. It provides the framework for agents to negotiate over resources, coordinate their tasks, and cooperate toward shared goals. Without this ability to exchange information and intent, individual agents would be unable to manage the complexities of operating in a shared environment with other goal-directed actors effectively.

107. (A) how well a candidate solution solves the problem.
The fitness function in a Genetic Algorithm is a quantitative measure that evaluates the performance of each individual in the population. It assigns a score based on how effectively a specific candidate solution satisfies the problem's objectives. This score is then used by the selection operator to determine which individuals are most likely to survive and reproduce in the next generation.

108. (A) 1, 3, 4, 2
A fuzzy inference system follows a logical four-step process. It begins with fuzzification, where crisp inputs are converted into fuzzy sets. Next, rule evaluation determines the influence of various fuzzy if-then rules. These results are then combined during the aggregation phase. Finally, the aggregated fuzzy output is converted back into a crisp value through defuzzification to produce a specific physical action.

109. (C) The exponential growth of the search space it needs to explore.
The Min-Max algorithm's primary drawback is the computational cost associated with exploring the entire game tree, which grows exponentially with search depth. Alpha-Beta pruning addresses this limitation by logically eliminating branches that cannot affect the final outcome. By reducing the number of nodes evaluated, it allows the search to reach significantly greater depths within the same time and memory constraints efficiently.

110. (B) the problem is about representing declarative knowledge with complex relationships.
State space representations are designed for problems that can be modeled as transitions between discrete states, such as pathfinding or games. However, they are not well-suited for representing vast amounts of declarative knowledge or complex, interconnected relationships between concepts. Such tasks are better handled by knowledge representation formalisms like logic, semantic networks, or ontologies, which focus on semantics and complex reasoning.

111. (A) context and intent.
Pragmatics is the level of linguistic analysis that focuses on how language is used in actual social situations. It considers the speaker's intent and the surrounding context to interpret meaning beyond literal semantics. For example, identifying Can you open the window? as a request rather than a literal inquiry about physical strength requires pragmatic reasoning about social norms and immediate circumstances.

112. (B) P-2, Q-3, R-4, S-1
Artificial intelligence involves diverse concepts with specific purposes. STRIPS is a formal language for planning operators. Hopfield networks provide a model for content-addressable memory. Conceptual Dependency theory offers a framework for representing the semantic meaning of sentences. Rational agents are models of entities that take actions to maximize their performance and achieve the best possible outcomes based on their perceived information.

113. (A) Hierarchical planning involves refining abstract plans into concrete steps.
Hierarchical planning manages complex problems by first creating an abstract plan and then refining each high-level step into more specific actions. Other planning types have different characteristics: Goal Stack Planning is fundamentally linear, and Partial Order Planning avoids committing to a total order early on. Linear planning is particularly susceptible to subgoal interactions, which hierarchical and non-linear planners are designed to solve.

114. (B) the principles of natural evolution and survival of the fittest.

representation helps the agent make more informed decisions by considering historical information and unobserved changes.

167. (C) Frames with inheritance
Frame-based systems with inheritance are ideal for representing general knowledge and its exceptions. A general Bird frame can contain the default attribute can-fly: yes. A more specific Penguin frame can inherit properties from the Bird frame but override the specific flying attribute with can-fly: no. This hierarchical structure efficiently manages commonalities and differences within a knowledge domain.

168. (A) Action A achieves a precondition P required by Action B.
In Partial Order Planning, a causal link is a formal way to track dependencies between actions. It records that a specific action was added to the plan for the purpose of fulfilling a precondition required by another action. By maintaining these links, the planner can identify and resolve potential threats where other actions might negate those essential conditions during execution.

169. (B) The goals of the agents are in conflict.
Competitive multi-agent systems are defined by the conflicting nature of individual agent objectives. In these environments, the success of one agent often comes at the expense of others, typical of games or market resource allocation. Agents must prioritize their own utility and anticipate the strategic moves of their adversaries, rather than collaborating toward a single, unified global goal.

170. (A) Both A and R are true and R is the correct explanation of A.
The selection operator in a Genetic Algorithm implements the principle of survival of the fittest. By assigning a higher probability of reproduction to individuals with higher fitness scores, the algorithm ensures that successful traits are passed on to the next generation. This biased selection process effectively guides the population through the search space toward increasingly optimal solutions over time.

171. (A) feedforward neural network.
A Multilayer Perceptron is a classic example of a feedforward artificial neural network. In this architecture, information moves in only one direction—from the input layer, through one or more hidden layers, to the output layer. Unlike recurrent networks, MLPs do not contain cycles or loops, making them suitable for static pattern recognition and function approximation tasks.

172. (A) Morphological Analysis
Identifying word stems or roots, a process often involving stemming or lemmatization, is a core component of morphological analysis in Natural Language Processing. This stage examines the internal structure of words to understand how they are formed from smaller units called morphemes. By reducing different forms of a word to a common base, the system can more effectively process language.

NODIA

173. (D) Heuristic Search - A search that is guaranteed to be optimal but not complete.
AI concepts are often defined by their specific guarantees. The Turing Test, rational agents, and state spaces are correctly described by their standard definitions. However, heuristic search is not inherently guaranteed to be optimal; its optimality and completeness depend entirely on the specific algorithm used and the mathematical properties, such as admissibility and consistency, of the chosen heuristic function.

174. (A) It learns to map high-dimensional input data onto a lower-dimensional grid.
A Self-Organizing Map is a type of unsupervised neural network that excels at dimensionality reduction. Its primary function is to take complex, high-dimensional input data and project it onto a simpler, typically two-dimensional grid. This mapping preserves the topological relationships of the data, meaning that points that are close in the high-dimensional space remain close on the lower-dimensional map.

175. (A) A and B only
Fuzzy sets are defined by membership functions with values between zero and one, and crossover operators combine parent solutions in genetic algorithms. While reflex agents use current percepts, they are not necessarily simple; model-based ones track history. Furthermore, the Min-Max algorithm is specifically designed

for two-player adversarial games, rather than for single-agent search problems as incorrectly suggested.

176. (C) Partial Order Planning
Partial Order Planning is a non-linear planning technique that explicitly identifies and manages conflicts between actions. These conflicts, known as threats, occur when an action's effects negate a precondition required by another action. The POP algorithm resolves these threats by adding ordering constraints, ensuring the integrity of the causal links that connect actions within the developing plan.

177. (B) the language, protocol, and content of messages exchanged between agents.
The communicative aspect of a multi-agent system is focused on the social interaction between independent entities. It defines the formal structure of their interaction, including the shared languages for expression and the protocols for message exchange. This framework allows agents to share their internal states and intentions, which is necessary for coordinating complex tasks and negotiating over shared resources.

178. (A) Defuzzification
Defuzzification is the concluding stage in a fuzzy logic control system. After the inference engine processes fuzzy rules to produce a combined fuzzy output, this output must be translated back into a form that can be used in the physical world. Defuzzification converts the theoretical fuzzy result into a single, crisp numerical value that provides a specific control command for actuators.

179. (A) The algorithm's search will become essentially random, and it may not converge.
The mutation rate in a Genetic Algorithm must be carefully balanced. While some mutation is necessary to maintain genetic diversity, an excessively high rate will constantly disrupt the successful traits that the algorithm has evolved. This turns the process into a random search, preventing the population from converging toward an optimal solution and undermining the effectiveness of the selection process.

180. (A) gradient descent applied to neural networks.
Backpropagation is a fundamental method for training artificial neural networks through supervised learning. It is essentially an efficient way to apply gradient descent to the network's weights. By using the chain rule to calculate how the error changes as each weight is adjusted, the algorithm allows the model to iteratively update its parameters to minimize prediction errors over time.

181. (D) Acting Rationally (Rational Agent)
The acting rationally approach to AI focuses on creating systems that behave as rational agents. A rational agent perceives its environment and selects the action that is most likely to maximize its expected performance measure. This approach is more general than human imitation, as it defines intelligence in terms of achieving optimal outcomes according to objective, mathematically analyzable criteria.

182. (B) Applies the rules to the known facts to deduce new information.
The inference engine is the reasoning component of an expert system. It acts upon the specialized information stored in the knowledge base by applying logical rules to current data. Through processes like forward or backward chaining, the engine deduces new facts or reaches conclusions, effectively simulating the analytical and diagnostic processes used by a human domain expert.

183. (B) Semantics is about the literal meaning of words/sentences, while pragmatics is about meaning in context.
The distinction between semantics and pragmatics is fundamental to Natural Language Processing. Semantics focuses on the objective, literal meaning of linguistic expressions independent of situational factors. Pragmatics, however, studies how the context, the speaker's intent, and social norms influence meaning. Understanding both levels is necessary for an AI to correctly interpret the full nuance of human communication.

184. (B) P-3, Q-1, R-4, S-2
Planning in AI involves specific concepts and formalisms. The Sussman Anomaly is a classic failure of linear planning. Hierarchical planning simplifies tasks using abstraction levels. Causal

Downloaded from www.pdf.tube

links in Partial Order Planning track how one action enables another. The STRIPS formalism provides a standard way to represent actions using three specific lists: preconditions, an add list, and a delete list for operators.

185. (A) The agent can operate without direct intervention from humans or other agents.
Autonomy is a core characteristic of an agent in a multi-agent system. It means that the agent possesses control over its own internal state and actions. Once deployed, an autonomous agent can perceive its environment and make decisions to achieve its objectives without requiring continuous external guidance or direct intervention from human operators or other automated entities.

186. (D) Encoding
Encoding is a critical preparatory step in Genetic Algorithms that defines how candidate solutions are represented. Before the algorithm can apply evolutionary operators, the problem's solution space must be mapped into a format the GA can process, such as bit strings or real-valued vectors. This chosen representation, or encoding, directly influences the effectiveness of crossover and mutation operators.

187. (B) It is a recurrent network that can act as an associative memory.
A Hopfield network is a type of recurrent artificial neural network, meaning it has internal feedback loops. This architecture allows it to function as an associative or content-addressable memory system. It can store multiple patterns and, when presented with a noisy or partial input, converge toward the closest stored pattern, effectively retrieving the original information from incomplete input cues.

188. (A) Reinforcement Learning
Learning to play a complex game like chess based solely on ultimate win or loss signals is a classic application of reinforcement learning. In this paradigm, an agent makes a long series of decisions and only receives feedback occasionally. Through trial and error, the agent learns to associate specific sequences of moves with positive or negative outcomes, eventually developing an optimal policy.

NODIA

189. (D) Alpha-Beta pruning is an optimization that returns the same result as Min-Max but by examining fewer nodes.
Alpha-beta pruning is strictly an optimization of the minimax search algorithm. It uses logical pruning to eliminate subtrees that are guaranteed not to affect the final move choice. Importantly, alpha-beta pruning is guaranteed to return the exact same move as a full minimax search to the same depth, but it does so more efficiently by evaluating significantly fewer nodes.

190. (C) A, B, C, and D only
Knowledge representation utilizes diverse models like ontologies for formal specifications, semantic networks for graphical relationships, and scripts for sequential scenarios. Conceptual Dependency provides primitives for meaning. While frame systems efficiently represent attributes through slots and fillers, they are primarily designed for declarative, object-oriented knowledge rather than being inherently specialized for representing procedural knowledge compared to other specific AI formalisms.

191. (B) It can get stuck in loops or undo previously achieved goals (e.g., Sussman Anomaly).
Goal Stack Planning is a linear technique that processes subgoals in a strict sequence. Its main limitation is the inability to effectively handle interdependent subgoals. This often results in the system undoing progress on one goal to satisfy another, a failure known as the Sussman Anomaly. This fundamental issue led to the development of more flexible non-linear planning methods.

192. (A) An agent program.
The agent program is the internal logic or implementation of an intelligent agent. It is the specific function that takes the current percept sequence as input and produces a corresponding action as output. While the agent architecture provides the machinery for perception and action, the agent program embodies the decision-making rules that determine the agent's actual behavior.

193. (A) membership function.
A fuzzy set is defined by its membership function, which assigns a grade of membership to each element in the universe of discourse.

The Min-Max algorithm is based on a conservative, game-theoretic assumption. It assumes that the opponent is a rational player who will always choose the move that is best for them. In a zero-sum game, this means the opponent will play optimally to minimize the first player's score, requiring the algorithm to find a move that is robust against perfect play.

230. (B) a small set of conceptual primitives representing actions, states, and objects.
Conceptual Dependency theory aims to represent the underlying semantic meaning of sentences in a language-independent way. It reduces complex linguistic expressions to a small, fixed set of fundamental primitives that describe basic actions and states. By focusing on these core conceptual components, the theory allows for sophisticated reasoning and inference that is independent of specific grammatical or lexical choices.

231. (B) manage complexity by planning at different levels of abstraction.
The main goal of hierarchical planning is to address the overwhelming complexity of large-scale problems. By breaking a task into multiple levels of abstraction, the planner can focus on the big picture first. This approach significantly reduces the initial search space, making it much more efficient to refine high-level goals into specific, actionable steps than searching at the lowest level.

232. (C) Pragmatics
Pragmatics is the branch of linguistics and Natural Language Processing that studies how context influences the use and interpretation of language. It examines how people use language to achieve specific social goals and how situational factors, such as the speaker's identity or environment, affect meaning. Understanding pragmatics is essential for AI to interpret human intent and social nuances correctly.

233. (A) percepts to actions.
An agent architecture is essentially the structural blueprint of an intelligent entity. It defines the mapping between the inputs the agent receives from its sensors, known as percepts, and the outputs it produces through its actuators, known as actions. This internal design determines the agent's behavior and its capability to process environmental information to successfully achieve its objectives.

234. (B) maximum of their membership functions.
In fuzzy set theory, standard operations like union and intersection are redefined to handle degrees of membership. For the union of two fuzzy sets, the most common method is the Zadeh operator, which takes the maximum value of the two membership functions for each element. This reflects the logical OR operation, where an element's membership in the union is determined by its strongest belonging.

235. (A) Selection, Crossover, and Mutation
A simple Genetic Algorithm relies on a cycle of three fundamental operators to drive the evolutionary process. Selection chooses the fittest individuals from the population to act as parents. Crossover then combines the genetic material of these parents to produce new offspring. Finally, mutation introduces random changes to maintain genetic diversity. These operators together allow the population to evolve toward an optimal solution.

236. (A) The chain rule for derivatives
The backpropagation algorithm is fundamentally a clever application of the chain rule from calculus. It allows the network to calculate how the overall error changes with respect to each individual weight across multiple layers. By systematically computing these partial derivatives from the output backward toward the input, the algorithm provides the necessary information to adjust weights and minimize prediction error during training.

237. (C) visualization and exploratory data analysis.
Self-Organizing Maps are highly effective tools for exploring complex data. By mapping high-dimensional input patterns onto a two-dimensional grid while preserving topological relationships, they provide a powerful visual representation of data clusters and trends. This makes SOMs ideal for exploratory data analysis, as they allow researchers to see similarities and relationships in large datasets that are otherwise difficult to perceive easily.

Click Here to Buy book on Amazon

search. Crucially, it only eliminates branches that it can mathematically prove will not influence the final decision. Therefore, the move selected by a minimax search with alpha-beta pruning is always identical to the move that would have been chosen by a full, exhaustive minimax search.

203. (B) determine the grammatical structure of a sentence.
The primary objective of parsing in Natural Language Processing is to uncover the underlying grammatical structure of a sentence. By applying formal grammar rules to a sequence of tokens, the parser identifies how words relate to each other as subjects, objects, or modifiers. This structural analysis is a necessary foundation for higher-level tasks like semantic interpretation and translation.

204. (C) A multi-agent system consists of a single agent solving a complex task.
By definition, a multi-agent system involves multiple autonomous entities interacting within a shared environment. These agents can be cooperative, competitive, or a mix of both. Suggesting that such a system consists of only a single agent is a contradiction. The power of multi-agent systems lies in the collective behavior and interaction of multiple goal-directed actors to solve complex problems.

205. (A) greater than 0.
The support of a fuzzy set is a fundamental property that identifies which elements of the universe have at least some level of membership. Specifically, it is the crisp set containing all elements whose membership value in the fuzzy set is strictly greater than zero. Elements with a membership value of exactly zero are considered outside the support of the set.

206. (C) both exploration and exploitation, by combining existing good solutions.
The crossover operator in Genetic Algorithms serves a dual purpose. It exploits known good regions of the search space by combining successful traits from two high-fitness parent solutions. Simultaneously, it facilitates exploration by creating new, untested genetic combinations in the offspring. This balance is essential for the algorithm to evolve superior solutions while maintaining the ability to discover better areas.

207. (B) hidden layers.
A multilayer perceptron is characterized by its hierarchical structure. It starts with an input layer and ends with an output layer. Between these are one or more hidden layers consisting of neurons with non-linear activation functions. These hidden layers are the engine of the network, enabling it to learn complex feature representations and solve problems that are not linearly separable.

208. (A) availability of labeled training data.
The primary distinction between supervised and unsupervised learning is the type of data used for training. Supervised learning relies on a dataset where each input is associated with a specific, correct label or output. In contrast, unsupervised learning processes data without any pre-defined labels, aiming instead to discover the inherent structure, clusters, or patterns within the input space itself.

209. (A) 3, 2, 1
Planning systems operate at various levels of specificity. The process begins with high-level goals that define the desired final outcome. These goals are then decomposed into more specific abstract plans, which outline the general strategy needed. Finally, these abstract steps are refined into concrete, actionable operations that can be directly executed by an agent to modify the environmental state.

210. (B) Multi-Agent Systems.
The field of multi-agent systems is centered on the study of autonomous, goal-directed entities. These agents are characterized by their ability to be proactive in pursuing objectives, reactive to environmental changes, and social in their interactions with other agents. Understanding how these entities coordinate, compete, and communicate in a shared environment is the primary focus of this specific area of AI.

211. (B) Emotional Intelligence approach
Standard approaches to AI are typically categorized by whether they focus on human-like behavior, human-like thinking, rational behavior, or rational thinking. While emotional intelligence

NODIA

This function is the core of fuzzy logic, as it maps inputs to values between zero and one. These values represent the degree to which an element belongs to the set, determining its role in rule inference.

194. (A) the best individual(s) from the current generation are guaranteed to be carried over to the next.
Elitism is a common refinement in Genetic Algorithms designed to preserve the best discovered solutions. Without elitism, the stochastic nature of crossover and mutation might inadvertently destroy the highest-fitness individuals. By guaranteeing that the top-performing solutions are carried over directly to the next generation, elitism ensures that the overall quality of the population never decreases over generations.

195. (A) Both Statement I and Statement II are true, and Statement II is the correct explanation of Statement I.
Training a multi-layer neural network involves optimizing weights over an error surface. Because these surfaces are typically non-convex, they contain many local minima where the error is lower than in the immediate surrounding area but higher than the global optimum. Consequently, the gradient descent-based backpropagation algorithm can become trapped in these local points, failing to find the overall best solution.

196. (B) Acting Humanly
The Turing Test is the classic benchmark for the acting humanly approach to AI. It evaluates whether a machine can behave in a way that is indistinguishable from a human in a natural language conversation. This approach focuses on external behavior and performance relative to a human baseline, rather than internal thought processes or the mathematical rationality of decisions.

197. (C) Heuristic value
In the STRIPS planning formalism, an action or operator is defined by its preconditions, an add list, and a delete list. These components provide a logical description of what must be true to execute an action and how that action changes the world state. While heuristic values are vital for guiding the search for a plan, they are not part of the operator's definition.

198. (A) Competitive
The interaction between independent, profit-seeking ride-hailing services in the same market is fundamentally competitive. Each service operates as an autonomous agent trying to maximize its own utility, often at the expense of its rivals. In this environment, agents compete for limited resources—passengers—making their goals inherently conflicting and requiring strategic decision-making to succeed in the shared environment.

199. (A) Fitness function
The fitness function is the evaluation component of a Genetic Algorithm. It takes a candidate solution and returns a numerical score reflecting how well it solves the problem. This measure is essential because it informs the selection process, allowing the algorithm to prioritize higher-quality solutions for reproduction and drive the population toward the optimal or near-optimal solutions.

200. (B) Unsupervised
Self-Organizing Maps are a prominent class of neural networks that utilize unsupervised learning. They do not require labeled data or a teacher signal; instead, they learn to represent the underlying structure and topology of the input space through competitive learning. This allows the SOM to cluster similar data points together on a low-dimensional grid for future analysis and visualization.

201. (C) Frames
Frame-based representation is a structured knowledge representation method where concepts are organized into entities called frames. Each frame contains multiple slots, which are attributes or properties of the concept. These slots are filled with specific values or pointers to other frames, creating a hierarchical and networked system of information that is highly effective for representing stereotypical objects and situations.

202. (A) Both A and R are true and R is the correct explanation of A.
Alpha-beta pruning is a logical optimization that significantly increases the efficiency of game tree

Downloaded from www.pdf.tube

any labeled examples or external teacher signals. By creating a low-dimensional map of the input space, SOMs provide a simplified representation that clusters similar data points, making them ideal for data visualization and exploration.

221. (D) Semantic Networks
Semantic networks are a graphical knowledge representation scheme where nodes represent objects or concepts and labeled arcs represent the semantic relationships between them. This structure is highly intuitive for modeling interconnected information and allows for efficient reasoning through property inheritance. By following paths between nodes, an AI can infer relationships and facts implied by the structure of the network itself.

222. (C) To create a formal, explicit specification of the concepts and relationships in a domain.
An ontology provides a structured and standardized way to represent knowledge within a specific domain. By formally defining the key concepts, their properties, and the relationships between them, it establishes a common vocabulary for agents and systems. This explicit specification is essential for ensuring that information is shared, integrated, and understood consistently across diverse applications and various organizational boundaries.

223. (A) Competitive
In a multi-agent system where independent entities bid for limited resources, the interaction is fundamentally competitive. Each agent attempts to maximize its own utility by securing resources at the best possible price. Since one agent's successful bid prevents others from acquiring the same resource, the goals of the participating agents are in direct conflict, requiring strategic negotiation and decision-making.

224. (B) A set of fuzzy if-then rules and membership functions
A fuzzy rule-based system is built upon two essential core components. Membership functions define the degrees to which crisp inputs belong to various fuzzy sets, handling the imprecision of data. Fuzzy if-then rules then provide the logic to map these fuzzy inputs to corresponding fuzzy outputs. Together, they allow the system to simulate human-like reasoning processes in uncertain environments.

225. (B) B, C, and D only
Modern artificial neural networks include several important models and algorithms. Backpropagation is the standard supervised training method, Hopfield networks are recurrent structures for memory, and Self-Organizing Maps provide low-dimensional representations of data. However, a single perceptron is fundamentally limited to solving linearly separable problems and therefore cannot learn the XOR function, which requires a multilayer architecture with hidden layers.

226. (C) guides the search by evaluating the quality of solutions.
The fitness function is the primary mechanism that guides the evolutionary process in a Genetic Algorithm. By evaluating each candidate solution and assigning it a numerical score based on its quality, the function provides the feedback necessary for selection. This ensures that the algorithm prioritizes the most successful individuals for reproduction, effectively steering the population toward increasingly optimal solutions over time.

227. (A) output layer back to the input layer.
The name backpropagation describes the direction in which the algorithm calculates the error gradients. After a forward pass produces an output, the system compares it to the target to find the error. This error information is then sent backward from the output layer, through any hidden layers, to the input layer, allowing for the systematic adjustment of network weights during training.

NODIA

228. (C) Training a robot to navigate a maze by rewarding it for moving closer to the exit.
Reinforcement learning involves an agent learning through interaction and environmental feedback. Training a robot to navigate a maze using rewards for progress is a classic example. The robot does not have a labeled dataset or a set of pre-defined patterns; instead, it must discover the optimal sequence of actions through trial and error based on the scalar rewards received.

229. (C) plays optimally to maximize their own score.

is an important area of research within AI, it is not considered one of the four fundamental, defining approaches to the field, which prioritize objective performance measures and logical or behavioral indistinguishability from humans.

212. (A) P-4, Q-3, R-2, S-1
AI involves several specialized techniques. An ontology is a formal specification of knowledge in a domain. STRIPS is a formalism used to represent planning operators. Backpropagation is the standard algorithm for training multilayer neural networks by minimizing error. Alpha-beta is a critical pruning technique used to optimize search in game trees by eliminating irrelevant branches, thus improving overall computational efficiency.

213. (D) First-Order Predicate Logic
First-order predicate logic is a powerful tool for representing knowledge, but it is fundamentally designed for precise, binary truth. It lacks the built-in mechanisms required to handle the degrees of belief or probabilistic reasoning found in modern AI. Conversely, methods like fuzzy logic, Bayesian networks, and certainty factors are explicitly created to manage the ambiguity and imprecision inherent in real-world information.

214. (A) Non-linear planning addresses subgoal interactions by interleaving plans.
Non-linear planning is a flexible approach that recognizes dependencies between different subgoals. Unlike linear planning, which solves subgoals in a strict sequence and often fails when they interact, non-linear planners can interleave actions from different parts of a plan. This capability allows the system to resolve conflicts and threats, ensuring that all goals are achieved efficiently without undoing previous progress.

215. (A) A, B, C, and D only
Genetic algorithms rely on several key components: encoding for representation, problem-specific fitness functions for evaluation, and operators like crossover and mutation to evolve solutions. Crucially, a Genetic Algorithm operates on a population of candidate solutions simultaneously, rather than a single solution. This population-based approach is what enables the algorithm to search diverse areas of the potential solution space effectively.

216. (C) content-addressable memory.
A Hopfield network is a recurrent neural network whose dynamics lead toward stable states, which correspond to local minima of an energy function. This property makes it ideal for content-addressable memory. When the network is initialized with a partial or noisy version of a stored pattern, it naturally evolves toward that stable state, effectively retrieving the complete and original information accurately.

217. (B) percept sequence, knowledge, and goals.
The behavior of a rational agent is the result of a complex decision-making process. It considers the entire sequence of percepts received from the environment, its internal knowledge base about the world, and its specific goals or utility functions. By integrating this information, the agent selects the action that maximizes its expected performance measure, demonstrating intelligent, goal-oriented behavioral patterns.

218. (C) Parsing
Parsing is the specific Natural Language Processing task of analyzing a sentence to determine its structural organization. By applying grammar rules, a parser identifies the hierarchical relationships between words and phrases, typically producing a parse tree. This process is essential for understanding how various parts of a sentence combine to convey meaning, serving as a bridge between tokens and semantics.

219. (C) Simple Linear Planning (like Goal Stack Planning)
The Sussman Anomaly is a classic problem that exposes the fundamental weakness of strictly linear planning. In these systems, subgoals are solved one after another in a fixed order. This often leads to failures when the actions needed for one subgoal necessarily undo progress made on another. This limitation was the primary motivation for developing more flexible and powerful non-linear planners.

220. (A) Both A and R are true and R is the correct explanation of A.
Self-Organizing Maps are a powerful example of unsupervised learning in neural networks. They function by discovering the inherent patterns and similarities within a dataset without requiring

Click Here to Buy book on Amazon

238. (A) Both Statement I and Statement II are true, and Statement II is a direct consequence of Statement I.
In search algorithms, admissibility is a formal property of a heuristic function that ensures it never overestimates the actual cost to reach the goal. This optimism is crucial because it guarantees that the search will never overlook a shorter path. Consequently, if a heuristic is admissible, algorithms like A* are mathematically guaranteed to find the absolute optimal solution to the search problem.

239. (C) Script
A script is a specialized knowledge representation structure designed to model stereotypical sequences of actions and events. It organizes information into temporal scenes, detailing the typical roles and props involved in common situations. This structured approach is particularly useful for tasks like story understanding, as it allows an AI to infer what happens next based on a familiar social scenario.

240. (B) POP commits to ordering decisions only when necessary to resolve conflicts.
The primary distinction between Partial Order Planning and Total Order Planning is the level of commitment to action sequencing. While TOP forces a strict, total order from the beginning, POP follows the principle of least commitment. It only imposes ordering constraints when required to resolve threats or meet preconditions, maintaining the flexibility to explore different paths through the plan space more efficiently.

241. (C) Deductive Learning
Neural networks typically employ various learning paradigms such as supervised, unsupervised, and reinforcement learning. These methods focus on training the model using data, rewards, or discovering patterns. Deductive learning, however, is a symbolic AI approach rooted in formal logic, where conclusions are derived from known premises and axioms. It is not considered a standard type of learning within the connectionist neural framework.

242. (A) MAX, MIN
In the alpha-beta pruning algorithm, the two variables serve as bounds for the players' scores. Alpha represents the best, or highest, score that the maximizing player is guaranteed to achieve. Beta represents the best, or lowest, score that the minimizing player is assured of. These values are used to prune irrelevant branches that neither rational player would ever choose to take.

243. (A) Semantic analysis
Semantic analysis is the phase of Natural Language Processing dedicated to identifying the actual meaning of a sentence. While syntactic analysis looks at structural correctness, semantics focuses on the relationships between words and the concepts they denote. This process is essential for understanding the information being conveyed, enabling machines to perform reasoning, answer questions, and meaningfully process human linguistic communication.

244. (C) A chromosome
In a Genetic Algorithm, the fundamental unit of representation is the chromosome. A chromosome is an encoded string of values, such as bits or real numbers, that represents a single candidate solution to the problem being solved. Each chromosome's fitness is evaluated, and the most successful ones are selected to pass their genetic material on to the next generation of solutions.

245. (B) linearly separable.
A single-layer perceptron is a basic neural network model that acts as a linear classifier. It can only successfully solve problems where the different classes can be divided by a single straight line, plane, or hyperplane. This limitation means it cannot learn more complex, non-linear functions like XOR unless it is expanded into a multilayer architecture with intermediate hidden layers and non-linearities.

246. (C) convert the final fuzzy output set into a single crisp value.
Defuzzification is the final stage of a fuzzy inference system. After applying logical rules to fuzzy inputs, the system generates a combined fuzzy set as its output. To be useful in the real world, this fuzzy result must be translated back into a specific, crisp numerical value. This value provides a clear command or decision that can be executed directly by a mechanical actuator.

247. (A) autonomy and proactive behavior.

Downloaded from www.pdf.tube

While both agents and objects encapsulate information and functionality, their level of initiative differs. Objects are reactive, only performing tasks when their methods are called by another entity. In contrast, agents are autonomous and proactive; they can take independent initiative to achieve their goals and respond dynamically to environmental changes. This self-directed behavior is the key distinction of agency in intelligent software systems.

248. (A) Preconditions, Add list, and Delete list.
The STRIPS planning model utilizes a specific three-part formalism to define actions. Preconditions list the facts that must be true for the action to occur. The add list identifies the new facts that will become true after the action, while the delete list specifies those that will become false. This structure allows the planner to efficiently track state changes and find valid paths to goals.

249. (C) Self-Organizing Map (SOM)
Self-organizing maps are a classic example of neural networks that learn without labeled data. Unlike multilayer perceptrons or backpropagation networks which are supervised, a SOM identifies the underlying patterns and structure of the input space through competitive learning. This allows it to create a low-dimensional map that clusters similar information together, making it an excellent tool for data visualization and high-dimensional analysis.

250. (D) conversational indistinguishability from a human.
The core concept of the Turing Test is that intelligence can be defined by observable behavior rather than internal mechanisms. It proposes that if a machine can communicate so effectively that a human judge cannot tell it apart from a person, then the machine has demonstrated a level of intelligence equivalent to a human. This approach prioritizes linguistic competence as the ultimate benchmark of intelligence.

251. (B) The machine's indistinguishability from a human in a textual conversation.
The Turing Test determines intelligence by assessing if a machine's behavioral output is indistinguishable from a human's during textual interactions. A human interrogator poses written questions to both a machine and a human. If the interrogator cannot reliably differentiate between the two based on their responses, the machine is considered to have passed the test successfully.

252. (A) An entity that acts to maximize its performance measure based on available information.
A rational agent is defined by its ability to act in a way that yields the best possible outcome or the highest expected outcome given its perceptions and internal knowledge. This paradigm emphasizes performance metrics and objective results rather than human-like thinking. It uses sensors to perceive its environment and actuators to perform actions that optimize its specific goals.

253. (C) It is based on a well-defined and measurable concept of rationality.
The rational agent approach is considered scientifically more tractable because it relies on mathematically defined and measurable criteria for success. Unlike the subjective and ambiguous nature of mimicking human thought processes, rationality can be optimized through engineering. This focus on performance against a specific measure allows for formal analysis and scientific evaluation of an artificial system's capabilities.

254. (D) A mechanism for perception (sensors) and a mechanism for action (actuators).
According to the framework of intelligent agents, any entity in this category must possess a mechanism for perception and a mechanism for action. Sensors allow the agent to receive input from its surrounding environment, forming a percept sequence. Actuators enable the agent to interact with and affect that environment based on the decisions made by its internal processing.

NODIA

255. (C) The rational agent approach
Designing an AI to prioritize the most effective strategy for winning a game represents the rational agent paradigm. This approach focuses on achieving the best possible performance measure rather than attempting to replicate human cognitive styles or behaviors. By optimizing for the best expected outcome, the system functions as an agent that selects actions strictly to fulfill its objective goal.

256. (B) An initial state, a set of operators, and one or more goal states.

A formal state space representation requires three fundamental elements to define a problem-solving environment effectively. These components include an initial state to mark the starting point, a set of operators or actions that govern transitions between states, and one or more goal states representing the solution. This structure allows problems to be visualized and solved as a graph traversal task.

257. (D) They use problem-specific knowledge to reduce the search space and improve efficiency.
Heuristic search techniques are preferred for complex problems because they incorporate problem-specific information to guide the exploration process. While uninformed methods like Breadth-First Search exhaustively visit nodes, heuristic strategies use estimation to prioritize paths likely to reach the goal. This approach significantly reduces the size of the search space, making solutions computationally feasible in vast and difficult domains.

258. (A) To estimate the cost of reaching a goal from the current state 'n'.
The heuristic function provides an informed estimate of the remaining cost to reach a goal state from any given current state. By calculating this value, the algorithm can focus on exploring the most promising directions first. This guidance allows the search process to prune less efficient paths, enhancing the overall performance and speed compared to using simple uninformed search methods.

259. (C) The quality of the guiding heuristic function.
The overall efficiency and success of a heuristic search are primarily determined by the accuracy and quality of the heuristic function employed. A well-designed function provides better guidance, leading to more significant performance gains and effective pruning of the search space. If the estimate is poor, the algorithm may fail to find solutions quickly or become computationally expensive.

260. (A) A task of finding a path in a graph.
Framing a problem within a state space allows it to be conceptualized as finding a path through a graph from a starting node to a destination node. This mathematical transformation simplifies complex problem-solving into a traversal task. By defining states as nodes and actions as edges, various search algorithms can systematically explore the connections to locate a viable solution path.

261. (D) The opponent will also play optimally to achieve their best outcome.
The Min-Max algorithm operates on the fundamental assumption that both players in a competitive environment will make the best possible decisions to further their own interests. It presumes the opposing player will always choose moves that minimize the current player's score. This adversarial model allows the algorithm to calculate a reliable strategy by preparing for the opponent's most effective counter-moves.

262. (B) To select moves that lead to states with the maximum possible score.
In the Min-Max framework, the player designated as MAX aims to achieve the highest possible utility value by the end of the game. During the search process, this player evaluates child nodes and selects the branch that leads to the maximum score. This strategy assumes the opponent will attempt to minimize that same score, creating a balanced search for optimal moves.

263. (C) To reduce the number of nodes evaluated by the Min-Max algorithm.
Alpha-Beta pruning serves as an optimization technique designed to increase the efficiency of the Min-Max search algorithm. Its primary objective is to decrease the number of nodes that must be evaluated within the game tree. By identifying and skipping branches that cannot influence the final decision, the search can go deeper or finish much faster than the standard version.

264. (D) By ignoring branches that are guaranteed not to affect the final optimal decision.
Alpha-Beta pruning achieves optimization by maintaining two threshold values that track the best scores found so far for both players. When the search identifies a branch that is guaranteed to be worse than a previously examined path, it stops exploring that section. This systematic ignoring of irrelevant nodes allows the algorithm to reach the same final decision with significantly less computation.

Click Here to Buy book on Amazon

265. (A) High computational cost due to exhaustive exploration.
A significant limitation of the standard Min-Max algorithm is the extreme computational resources required to explore an entire game tree exhaustively. As the depth or branching factor of the game increases, the number of nodes to evaluate grows exponentially. This complexity often makes it impossible to search to the very end of many games, necessitating optimizations or depth limits.

266. (B) FOL can represent relationships between objects and make generalized statements.
First-Order Logic provides a significant advancement over Propositional Logic by allowing for the representation of complex relationships between distinct objects. It utilizes predicates and variables to describe properties and interactions, which Propositional Logic cannot handle. Additionally, the inclusion of quantifiers enables the system to make broad, generalized statements about entire sets of objects within a specific domain of knowledge.

267. (B) Nodes represent concepts, and arcs represent relationships between them.
In the graphical structure of a semantic network, nodes are used to represent individual concepts, objects, or classes of information. The directed arcs, or links, that connect these nodes signify the specific relationships or associations between them. This visual and structural approach allows for intuitive mapping of knowledge and facilitates reasoning through the connections established between different conceptual entities.

268. (A) Inheritance
The process where a specific concept automatically acquires the attributes and properties of a more general class is known as inheritance. In a semantic network, this occurs when a specialized node is connected to a general node via a specific link. This mechanism simplifies knowledge storage by allowing properties to be defined once at a high level and shared down the hierarchy.

269. (C) Quantifiers
First-Order Logic employs quantifiers as specialized constructs to formulate statements regarding the scope of objects in a domain. The universal quantifier is used to express facts that apply to every member of a set, while the existential quantifier indicates that a property holds for at least one member. These tools provide the necessary expressiveness for making general rules and observations.

270. (D) Semantic Network
A semantic network is inherently graphical because it formalizes knowledge using a structure composed of nodes and directed arcs. This distinguishes it from logical schemes like Propositional or First-Order Logic, which rely on symbolic and sentential formulas. The graph-based nature of semantic networks makes them particularly useful for visualizing hierarchical relationships and facilitating inheritance-based reasoning through connected conceptual paths.

271. (C) To hold a specific piece of information or attribute related to the concept.
In a frame-based representation system, a slot functions as a dedicated placeholder for a specific attribute or piece of information related to a stereotyped situation. Each frame consists of multiple slots that define the characteristics of the concept being modeled. These slots can be populated with specific values or default information, allowing the system to build a detailed description of objects.

272. (B) As a collection of 'IF-THEN' conditional statements.
Rule-based systems represent knowledge primarily through a collection of conditional statements known as IF-THEN rules. The first part of the rule specifies the conditions or premises that must be met, while the second part describes the resulting action or conclusion. This structure allows the system to store expertise in a modular format that is easily interpreted by an inference engine.

273. (D) Script
A script is a specialized knowledge representation structure designed to model a characteristic sequence of events for a specific context. It includes details such as entry conditions, expected roles, and a series of scenes that typically occur. This formalism allows artificial intelligence systems to make accurate inferences about the expected flow and outcome of familiar, recurring activities and social situations.

NODIA

274. (A) The action or conclusion specified in the rule's 'THEN' part is executed.
In the context of a rule-based system, a rule is said to fire when the current facts within the knowledge base satisfy the specific conditions defined in the rule's antecedent. Once triggered, the action or conclusion specified in the consequent part of the rule is executed. This process often results in new information being added to the system's internal knowledge base.

275. (C) Frames, as they allow for assumption-based reasoning.
Frame-based representations utilize default values within their slots to support assumption-based reasoning when specific data is missing. By providing a standard value for a property, such as a typical wall color, the system can make reasonable inferences about a situation. This feature enables artificial intelligence to maintain a complete model of an environment even when information is partially incomplete.

276. (A) To represent the meaning of sentences in a standardized, language-independent format.
The primary objective of Conceptual Dependency theory is to represent the underlying meaning of natural language sentences in a consistent and language-independent format. By breaking down linguistic input into a canonical structure, it ensures that different sentences expressing the same idea result in identical representations. This standardization simplifies the process of making semantic inferences across various language inputs.

277. (D) They serve as the fundamental building blocks for representing the meaning of actions.
Within Conceptual Dependency theory, primitive actions like PTRANS and ATRANS serve as the fundamental building blocks for representing the core meaning of various activities. PTRANS signifies a physical transfer of location, while ATRANS represents an abstract transfer of possession. These primitives, combined with specific conceptual roles, allow complex sentences to be decomposed into their basic, essential semantic components for processing.

278. (B) It is an explicit and formal specification of a conceptualization.
An ontology is defined by its role as an explicit and formal specification of a shared conceptualization within a particular domain. It provides a structured model that includes a standardized vocabulary of concepts, properties, and the relationships existing between them. This formality ensures that the knowledge is machine-interpretable and can be consistently understood by different systems, supporting rigorous reasoning processes.

279. (C) To enable knowledge sharing and interoperability between different systems.
A vital function of ontologies is to facilitate knowledge sharing and semantic interoperability among different artificial intelligence systems and agents. By providing a common, formalized framework for understanding a domain, ontologies allow disparate systems to integrate data and collaborate effectively. This capability is essential for the development of the Semantic Web and other large-scale, interconnected knowledge-based environments.

280. (D) higher degree of formality and use of standard languages.
While both represent relationships between concepts, ontologies are distinguished from simpler semantic networks by their much higher degree of formality and use of standardized languages like OWL. Ontologies are specifically engineered for rigorous machine-interpretable semantics and broad interoperability. This structured approach allows them to support more complex reasoning and data integration tasks across diverse and geographically distributed systems.

281. (A) A knowledge base and an inference engine.
The architecture of an expert system is built upon two essential components: a knowledge base and an inference engine. The knowledge base contains specialized, domain-specific facts and rules acquired from human experts. The inference engine acts as the reasoning mechanism that processes this information to solve problems. Together, these elements allow the system to emulate the decision-making capabilities of human experts.

282. (B) To apply knowledge to a problem to derive conclusions.

Downloaded from www.pdf.tube

representation that conveys the actual underlying meaning.

300. (A) Semantic analysis

The process of deriving the intended message of a sentence involves interpreting its syntactic structure and resolving various forms of ambiguity, such as word sense or pronoun references. This task, known as semantic analysis, focuses on understanding what the text actually communicates rather than just how it is constructed. It is the final stage in achieving deep comprehension within natural language processing systems.

www.nodia.in

NODIA

conditions relevant to the successful completion of the intended task or objective.

291. (C) The action is simulated, and the world state is updated.
In Goal Stack Planning, when an action is popped from the top of the stack and all its required preconditions are found to be satisfied, the planner simulates the execution of that action. This involves updating the current representation of the world state to reflect the effects of the action. Once the state is modified, the planning process continues with the next item.

292. (D) Preconditions, Add list, and Delete list.
The STRIPS representation formalizes actions using three distinct lists of literals. The preconditions define what must be true before the action can occur. The add list specifies the new facts that become true as a result of the action, while the delete list identifies the facts that are no longer true. This structure provides a clear and efficient way to model changes.

293. (A) It lists the new facts that become true after the action is executed.
Within the STRIPS formalism, the add list serves the specific function of enumerating all the new positive literals that will be added to the world state once an action is completed. It represents the constructive effects of the action, detailing the changes that make certain conditions true. This list is essential for the planner to track progress toward achieving goal literals.

294. (B) To create a plan at an abstract level first and then refine it with details.
Hierarchical Planning is based on the principle of generating a plan at a high level of abstraction before refining the details. The process begins by creating a simplified sequence using abstract operators, which is then progressively decomposed into more concrete, executable actions. This top-down strategy helps manage complexity by addressing the most significant aspects of the problem first before adding detail.

295. (C) Complex problems where abstracting away details can reduce search.
Hierarchical Planning is particularly advantageous for complex problems where many subgoals interact or the state space is very large. By first planning with abstract operators, the system can ignore minor details that would otherwise cause a combinatorial explosion in the search process. This reduction in complexity allows the planner to find viable high-level solutions more efficiently than non-hierarchical methods.

296. (D) To provide rules for forming syntactically correct sentences.
In Natural Language Processing, the primary role of a grammar is to provide a formal set of rules that dictate how words can be combined to create syntactically valid sentences. These rules define the structural relationships within the language, allowing the system to distinguish between correctly formed strings and nonsense. Grammar acts as the foundational framework for all syntactic analysis tasks.

297. (C) Parsing
Parsing is the specific process used to analyze the grammatical structure of a sentence based on the rules of a given grammar. During this phase, an algorithm examines the sequence of words to determine how they fit together into a hierarchical structure, such as a parse tree. This analysis is crucial for verifying that the input follows the language's formal syntactic requirements.

298. (A) Top-down starts from the sentence symbol, while bottom-up starts from the words.
The fundamental difference between top-down and bottom-up parsing lies in their starting points. Top-down parsing begins with the highest-level grammar symbol, such as the sentence indicator, and attempts to expand it to match the input words. Conversely, bottom-up parsing starts with the individual words of the sentence and tries to combine them into increasingly complex structures until the top-level symbol is reached.

299. (B) Because syntactic correctness does not guarantee that a sentence is meaningful.
Semantic analysis is a necessary step after parsing because establishing the syntactic correctness of a sentence does not guarantee its meaning is understood. While parsing identifies the grammatical roles and structure, it does not interpret the intended message or resolve ambiguities. Semantic analysis bridges this gap by mapping the syntactic structure to a formal

Click Here to Buy book on Amazon

The inference engine in an expert system serves as a general-purpose reasoning component that applies the rules stored in the knowledge base to the facts of a specific problem. Its main task is to systematically navigate through the available information to deduce new conclusions or recommend specific actions. This component effectively manages the logic required to provide expert-level advice or solutions.

283. (B) Because real-world problem domains often involve incomplete or uncertain information.
Handling uncertainty is a critical requirement for expert systems because real-world problem domains are frequently characterized by information that is incomplete, ambiguous, or inherently uncertain. Without a robust mechanism to manage these factors, the system would struggle to produce reliable conclusions. Incorporating uncertainty management allows the system to provide confidence levels and function effectively in complex, unpredictable practical environments.

284. (A) Techniques for managing and reasoning with uncertain information.
Certainty factors and Bayesian networks are prominent examples of techniques used by expert systems to manage and reason with uncertain or probabilistic information. Certainty factors assign numerical degrees of belief to rules and facts, while Bayesian networks use probability theory to propagate uncertainty through the system. Both methods enable the AI to reach conclusions that are qualified by statistical confidence.

285. (C) Human experts in a specific, narrow domain.
The information residing within the knowledge base of an expert system is typically obtained by extracting the specialized insights and experience of human experts. This knowledge is specifically focused on a narrow and well-defined domain of expertise. By encoding this human intelligence into a machine-readable format of facts and rules, the system can provide high-quality decision support in its area.

286. (D) To find a sequence of actions that leads from an initial state to a goal state.
AI planning is a deliberative process focused on identifying a specific sequence of actions that will successfully transition the world from an initial state to a desired goal state. This task requires a formalized representation of the environment, including the capabilities of available actions. The resulting plan serves as a roadmap for an agent to achieve its objectives through systematic behavior.

287. (C) The conditions that must be true for the action to be executable.
In the context of a classical planning system, the preconditions of an action represent the specific set of conditions or literals that must be true in the current world state for that action to be performed. If these requirements are not met, the operator cannot be executed. Defining preconditions is essential for determining which steps are valid during the plan generation process.

288. (B) It generates a plan as a strict, sequential ordering of actions.
A linear planner, also known as a totally ordered planner, is characterized by the construction of a plan as a single, strict sequence of consecutive steps. This approach determines the exact order of every action throughout the planning process. While simple, this method can be inefficient because it may force premature decisions about action order, potentially leading to unnecessary backtracking during search.

289. (D) It is more flexible as it does not commit to a specific action order unless necessary.
Non-linear planning offers increased flexibility by maintaining only a partial ordering of actions and avoiding specific sequence commitments until they are absolutely necessary to resolve conflicts. This allows the system to work on different subgoals independently, which can be more efficient than the rigid approach of linear planning. By postponing ordering constraints, it reduces the likelihood of reaching difficult dead ends.

290. (A) As a partial description of a desired state to be achieved.
In classical planning systems, a goal is typically represented as a partial description of the state the agent wishes to achieve. This is often formulated as a conjunction of specific literals that must become true in the final state. Unlike a complete world description, a goal only specifies the

NODIA

VISIT:

www.pdf.tube

Download PDFs for:

- ✓ Recommended Textbook Resources
- ✓ Mock Tests with Solutions
- ✓ Chapterwise Question Bank
- ✓ Previous Year UGC-NET Exam Papers
- ✓ Previous Year Assistant Professor Exam Papers

NODIA

Nodia Editions LLP

125, Sector 6, Vidhyadhar Nagar, Jaipur 302039

MRP. ₹ 710